PCT/GB2004/001349

INVESTOR IN PEOPLE

The Patent Office
Concept House
Cardiff Road
Newport
South Wales
NP10 8QQ

REC'D 04 MAY 2004
WIPO PCT

PRIORITY DOCUMENT

SUBMITTED OR TRANSMITTED IN COMPLIANCE WITH RULE 17.1(a) OR (b)

I, the undersigned, being an officer duly authorised in accordance with Section 74(1) and (4) of the Deregulation & Contracting Out Act 1994, to sign and issue certificates on behalf of the Comptroller-General, hereby certify that annexed hereto is a true copy of the documents as originally filed in connection with the patent application identified therein.

In accordance with the Patents (Companies Re-registration) Rules 1982, if a company named in this certificate and any accompanying documents has re-registered under the Companies Act 1980 with the same name as that with which it was registered immediately before re-registration save for the substitution as, or inclusion as, the last part of the name of the words "public limited company" or their equivalents in Welsh, references to the name of the company in this certificate and any accompanying documents shall be treated as references to the name with which it is so re-registered.

In accordance with the rules, the words "public limited company" may be replaced by p.l.c., plc, P.L.C. or PLC.

Re-registration under the Companies Act does not constitute a new legal entity but merely subjects the company to certain additional company law rules.

Signed M Cunki

Dated 27 April 2004

BEST AVAILABLE COPY

Patents Act 1977
(Rule 16)

THE PATENT OFFICE
31 MAR 2003
LONDON

01APR03 E796642-1 D02246
P01/7700 0.00-0307431.7

The Patent Office

Cardiff Road
Newport
South Wales
NP10 8QQ

Request for grant of a patent

(See the notes on the back of this form. You can also get an explanatory leaflet from the Patent Office to help you fill in this form)

1. Your reference	P016275GB
2. Patent application number *(The Patent Office will fill in this part)*	0307431.7
3. Full name, address and postcode of the or of each applicant *(underline all surnames)*	SONY UNITED KINGDOM LIMITED THE HEIGHTS BROOKLANDS WEYBRIDGE, KT13 0XW UNITED KINGDOM
Patents ADP number *(if you know it)*	652270E001
If the applicant is a corporate body, give the country/state of its incorporation	UNITED KINGDOM
4. Title of the invention	VIDEO PROCESSING
5. Name of your agent *(if you have one)*	D Young & Co
"Address for service" in the United Kingdom to which all correspondence should be sent *(including the postcode)*	21 New Fetter Lane London EC4A 1DA
Patents ADP number *(if you know it)*	59006

6. If you are declaring priority from one or more earlier patent applications, give the country and the date of filing of the or of each of these earlier applications and *(if you know it)* the or each application number	Country	Priority application number *(if you know it)*	Date of filing *(day / month / year)*

7. If this application is divided or otherwise derived from an earlier UK application, give the number and the filing date of the earlier application	Number of earlier application	Date of filing *(day / month / year)*

8. Is a statement of inventorship and of right to grant of a patent required in support of this request? *(Answer 'Yes' if:*
a) any applicant named in part 3 is not an inventor, or
b) there is an inventor who is not named as an applicant, or
c) any named applicant is a corporate body.
See note (d))

Yes

9. Enter the number of sheets for any of the following items you are filing with this form. Do not count copies of the same document

Continuation sheets of this form	0
Description	27
Claim(s)	4
Abstract	1
Drawing(s)	16 +16

10. If you are also filing any of the following, state how many against each item.

Priority documents	0
Translations of priority documents	0
Statement of inventorship and right to grant of a patent *(Patents Form 7/77)*	0
Request for preliminary examination and search *(Patents Form 9/77)*	1
Request for substantive examination *(Patents Form 10/77)*	0
Any other documents *(please specify)*	0

11. I/We request the grant of a patent on the basis of this application.

Signature D Young & Co Date 31 March 2003

D Young & Co (Agents for the Applicants)

12. Name and daytime telephone number of person to contact in the United Kingdom James Turner 023 8071 9500

Warning

After an application for a patent has been filed, the Comptroller of the Patent Office will consider whether publication or communication of the invention should be prohibited or restricted under Section 22 of the Patents Act 1977. You will be informed if it is necessary to prohibit or restrict your invention in this way. Furthermore, if you live in the United Kingdom, Section 23 of the Patents Act 1977 stops you from applying for a patent abroad without first getting written permission from the Patent Office unless an application has been filed at least 6 weeks beforehand in the United Kingdom for a patent for the same invention and either no direction prohibiting publication or communication has been given, or any such direction has been revoked.

Notes

a) *If you need help to fill in this form or you have any questions, please contact the Patent Office on 08459 500505.*

b) *Write your answers in capital letters using black ink or you may type them.*

c) *If there is not enough space for all the relevant details on any part of this form, please continue on a separate sheet of paper and write "see continuation sheet" in the relevant part(s). Any continuation sheet should be attached to this form.*

d) *If you have answered 'Yes' Patents Form 7/77 will need to be filed.*

e) *Once you have filled in the form you must remember to sign and date it.*

f) *For details of the fee and ways to pay please contact the Patent Office.*

VIDEO PROCESSING

The present invention relates to processing video data.

It is known to link video and audio devices in a television studio together using a switching device, typically a cross point switch.

The present inventors have identified the need for a system which links audio and video devices in a studio by a switched local area network, for example an Ethernet network, operating with a known protocol such as Internet Protocol (IP). However most, if not all, current video and audio equipment used in studios is not equipped to operate with such a switched network operating according to such a protocol.

The audio and video devices used in a studio include cameras, editors, audio mixers, video tape recorders (VTRs) and play-out switches amongst other examples. It is also known to use monitors to view video which is being played out or to preview on monitors video which is available to be played out.

It is desirable to provide a similar ability in a switched network.

According to one aspect of the present invention, there is provided a video processor connectable to a data network and being arranged to receive a digital video data stream of a first resolution, the processor comprising: a video generator arranged to produce from the video data stream of the first resolution a video data stream of a second, lower, resolution; a packetiser operable to format at least that part of the video data stream of the first resolution not represented by the video data stream of the second resolution into data packets and to format the video data stream of the second resolution into data packets; and a network interface arranged to launch the data packets onto the data network, the video generator, the packetiser and the network interface operating substantially in real time.

Thus for example, second, lower, resolution video data is produced from the first, higher, resolution video data. The higher and lower resolution video data are preferably launched onto a packet-based network in separate multicast groups so that, for example, the higher resolution video may be supplied to a first destination or set of destinations associated with the first multicast group and the lower resolution data may be supplied to a second, possibly different, destination or set of destinations associated with the second multicast group.

In this way, the higher resolution video supplied to the first destination(s) may be monitored in lower resolution at the second destination(s). Thus for example the processing

of the higher resolution video at the first destination may be controlled by control equipment at the second destination. The use of the lower resolution reduces the processing overhead and data traffic at such control equipment, potentially allowing several such lower resolution streams to be handled at the control equipment.

Preferably, to allow convenient selection of network recipients, the network interface is arranged to launch the data packets corresponding to the video data stream of the first resolution onto the network as a multicast group. Preferably the network interface is also arranged to launch the data packets corresponding to the video data stream of the second resolution onto the network as a second multicast group, different to the first multicast group. This allows different network recipients to be set up for the two streams.

Preferably, so that a low amount of processing is needed to display or handle the lower resolution data, the or each video data stream of the second resolution comprises an uncompressed video data stream. To provide a useful degree of compatibility with personal computer (PC) based equipment, it is preferred that the or each video data stream of the second resolution conforms with the RGB 555 format, representing each of red, green and blue pixel values by 5 bits of a pixel word. This format is particularly suited to allow display on a PC's display screen (particularly in a 16-bit colour mode of operation) with a relatively small amount of processing being required at the PC. Potentially, several streams at the low resolution could be displayed at the same time on a single PC's screen. For efficient transmission of such data, it is preferred that the packetiser is arranged to format the or each video data stream of the second resolution into RTP (Real time Transport Protocol) packets having a line of video data in each such packet.

The packetiser can be arranged to format the (whole of the) video data stream of the first resolution into data packets. This has the advantage that the video data stream of the first resolution is self-contained. Alternatively, where a subset of the video data stream of the first resolution can be derived from the respective video data stream of the second resolution, the packetiser can be arranged to format only a part, being all but the subset, of the video data stream of the first resolution into data packets. This has the advantage of reducing network data traffic, as unnecessary data is not transmitted.

As well as video data, audio and control data need to be transferred between source and destination equipment (possibly different equipment to video destination equipment). Advantageously, this can be carried out over the same network as the video data. To this end

it is preferred that the processor is arranged to receive one or more audio and/or control data streams associated with the digital video data stream; the packetiser is arranged to format the audio and/or control data streams into data packets; and the network interface is arranged to launch the data packets of the audio and/or control data streams onto the network.

In the case in which the video data stream of the first resolution represents an interlaced video signal, and in which the lower resolution video is produced from one field of the interlaced signal, a system operating in substantially real time could potentially suffer from relatively high network traffic in one field period and relatively low network traffic in the other field period. To allow the traffic generated by the lower resolution signal to be better balanced, it is preferred that the video generator is selectably operable to generate the video data stream of the second resolution from only odd fields or from only even fields of the video data stream of the first resolution. Alternatively, the video generator could be operable to generate a video data stream of the second resolution from only odd fields and a video data stream of the second resolution from only even fields of the video data stream of the first resolution. This allows network load balancing to be carried out by the network destination selecting which of the two (equally useful) lower resolution streams should be received from that video processor.

Although the processor can be a stand-alone unit connectable to video equipment, the processor is preferably embodied as a component of video source equipment (such as a camera, VTR, vision mixer etc) and/or as a component of video destination equipment (such as a monitor, VTR, vision mixer etc).

Further respective aspects and features of the present invention are defined in the appended claims.

Embodiments of the invention will now be described, by way of example only, with reference to the accompanying drawings in which:

Figure 1 is a schematic block diagram of a network in a studio;

Figure 2 is a schematic simplified diagram of the network showing data flows across the network;

Figure 3A is a schematic diagram of the format of an audio or video packet used in the network;

Figure 3B is a schematic diagram of the format of an AVSCP or CNMCP packet used in the network;

Figure 3C schematically illustrates a unicast data packet;

Figure 4 is a schematic block diagram of a network interface of the network of Figure1;

Figure 5A is a schematic diagram of the format of an data packet used in the network interface;

Figure 5B is a schematic example of current flow assignment;

Figures 5C schematically illustrates data flow in an ENIC;

Figures 6A and 6B schematically illustrate a packetiser/depacketiser switch of the network interface;

Figure 7 is a schematic block diagram of an illustrative small network for explaining a mode of operation of the network; and

Figure 8 is a schematic block diagram of a proxy generator of the network interface;

Figure 9 is a schematic diagram of one example of the display of a Graphical User Interface (GUI); and

Figure 10 is a schematic diagram of another example of the display of a Graphical User Interface (GUI);

Figure 11 is a schematic diagram of an example of a graphical interface for illustrating the configuration of the network;

Figure 12 is a schematic diagram of an example of a graphical interface for illustrating how data is routed across the network;

Figure 13 schematically illustrates a user interface provided on the network manager via which a user may enter configuration data;

Figure 14 schematically illustrates a protocol stack; and

Figure 15 schematically illustrates an AVSCP header.

Overview and Terminology

Referring to Figure 1, a network is installed in for example a studio. The network is a multicast network comprising an asynchronous *n*Gigabit Ethernet switch 2, where *n* is 1 or 10 for example. Connected to the switch 2 are source "groups" S1 to S10, destination "groups" D1, D2, D3, D8 and D9, and a network control arrangement, in this example a network manager 4 and one or more switching and routing clients 6, 61. In this context, "group" refers to a camera, VTR, DSP etc. which has one or more inputs and/or one or more outputs. Herein

inputs are termed "destination devices" and outputs are termed "source devices" with respect to video. As will become apparent, a destination group may also act as a source and a source group may also act as a destination. An example of a source group is a video tape recorder (VTR) that would have audio, video, status and proxy devices associated with it. The source groups S and destination groups D are coupled to the switch 2 by network interface cards NI 1 to 11 and which are also termed ENICs (Enhanced Network Interface Cards) herein. As indicated with respect to Cameras 1 and 2 and ENICs NI1 and 2, one group may be connected to two ENICs; i.e. one or more inputs and outputs of the group may be connected to one ENIC and the others to the other ENIC.

In a conventional studio, the source groups, e.g. cameras and destination groups e.g. video processors are connected by a cross point switch. The conventional cross point switch requires specific known devices to be connected to corresponding specific known ports on the switch to ensure that they can be connected together via switch. The network of Figure 1 including the switch 2 is configured by the network manager 4 and the switching and routing client 6 to emulate a cross point switch at least to the extent that any one or more sources can be connected to any one or more destinations. For that purpose the network is an IP multicast network using IGMP (Internet Group Management Protocol). Unlike the conventional cross point switch network, the network of Figure 1 connects source devices to destination devices according to identifiers identifying the source devices and destination devices and defining multicast addresses which the identified source devices and destination devices join so that the actual ports of the switch 2 to which they are connected are irrelevant.

It should be noted that this example of the network operates as follows: a single source device belongs to one and only one multicast group to which it transmits data. A destination device receives data from that source device by joining the source device's multicast group. That is done by issuing a multicast group join message. The invention is not limited to that: other ways of operating are possible.

Overview of ENICs

An ENIC is described in more detail in the section below headed ENIC. An ENIC allows any source group, for example a camera, and any destination group, for example a VTR, which is not designed for use with a multicast network to be used in a multicast network. An ENIC is a "dumb" device which can be requested to supply and receive audio, video, and control data streams. An ENIC cannot view or initiate any change to the

configuration of the network. An ENIC may subscribe to one or more multicast groups as directed by the network manager 4.

Each ENIC has an Ethernet address and an IP address. A multicast address and a UDP port are used to identify different data streams and in this embodiment a single IP address is associated with each ENIC. However, in alternative embodiments multiple IP addresses could be associated with a single ENIC. It also has an ENIC identifier ID and port Ids for respective ones of the destination devices and source devices of the ENIC. As will be described those addresses and IDs and other data are recorded with the network manager 4. The source devices and destination devices correspond to respective ones of one or more physical inputs and outputs of an ENIC. The source groups and destination groups may have "tally text" associated with them. An example of tally text is a name such as "VTR1" which may be given to a VTR or a cameraman's name e.g. Jim which may be given to a camera. The tally text is recorded at the network manager. All groups connected to the network may be named in this way. Source devices may have tally text, termed sub_tally herein, associated with them.

An ENIC acts as a switch which switches data received from the switch 2 to a specified physical output of the ENIC and switches data from a specified physical input to the switch 2.

The network, implemented using the switch 2, is asynchronous, but video and audio needs synchronous processing. The video and audio devices connected to the network operate on serial digital data, for example SDI. The ENICs convert serial digital video and audio to multicast UDP/IP packets and convert multicast UDP/IP video and audio packets to serial digital video and audio. The ENICs also provide synchronous operation over the network and, as will be described, align frames of different video streams for purposes such as editing. The ENICs also generate low resolution video streams referred to herein as "proxy video" as will be described.

The network optionally also comprises a master ENIC NIM which will be described in more detail in the section Frame Start Alignment below.

<u>Overview of Network Manager</u>

The network manager 4 which may comprise a computer, for example a Personal Computer (PC), is linked to the network via a standard network interface card. There may be only one network manager connected to the network. The network manager maintains a

database of the configuration of the network. It allocates resources to the client(s) 6 and ENICs NI, sends commands that change AV connections across the network and ensures the switching and routing client's view of the network is correct. The network manager, amongst other functions, defines respective multicast groups to which the source devices belong. In alternative embodiments a further network manager may be provided such that there is a master network manager and a slave network manager, which would provide a safeguard against failure of one of the network managers.

The manager records, in relation to each ENIC, the Ethernet address the IP address of the ENIC and a reference that is used as a reference in the network manager. The IDs of the source and destination devices of an ENIC are recorded at the network manager 4. The source groups and destination groups and may have "tally text" associated with them. An example of tally text is a name such as "VTR1" which may be given to a VTR or a cameraman's name e.g. Jim which may be given to a camera. The tally text is recorded at the network manager. In addition, a source device may have sub_tally text which is recorded at the manager 4. Other data is also recorded as described in the next section Network Configuration Data.

Network Configuration Data

The network manager stores and maintains a set of data relating to each of a number of different categories of device on the network.

In particular, the network configuration data is of four basic types relating to 4 different types of device.

1. SOURCE device:: Video, Audio and Status is prepared by an ENIC and transmitted to a multicast group on the network. Each SOURCE device can also transmit a proxy.
2. DESTINATION device:: Video, Audio and Status is received from the network by joining a multicast group.
3. CONTROL_SOURCE device: Control commands are generated by an ENIC or Client and are transmitted unicast to a CONTROL_DESTINATION.
4. CONTROL_DESTINATION device: Receives control commands unicast from a CONTROL_SOURCE.

The client 6 cannot directly access the SOURCE and CONTROL_DESTINATION devices. These devices are members of a CONTROL_SOURCE_GROUP, a structure that

groups devices that cannot be controlled independently. For example, SDI outputs from a VTR are connected to an ENIC for transmission onto the network 2. The SDI input is represented as two SOURCE devices *(Vo, A_0,)* in the network configuration. All of these devices are from the same physical device, are jointly controlled and have a common time code and stream status, i.e. stop, FF(fast forward), rew (rewind), etc.

A predetermined set of information is stored in relation to each of the above device types, for example a device identifier, a destination IP address associated with the device and a UDP port and/or IP address of an ENIC associated with the device.

A predetermined set of information associated with an ENIC known as an ENIC structure is held in a database in the network manager and comprises the Ethernet address and IP address of the ENIC. The ENIC structure also maps the associated devices (e.g. audio and video streams) to the physical ports on the card and includes any hardware limitations that restrict the ideal model described above. When an ENIC starts up it will receive information on what devices are connected to its RS422 ports (i.e. VTR, TALLY), so that the correct driver can be used.

Overview of Switching and Routing Client 6.

There may be one or more clients 6, 61 connected to the network but the following description assumes there is only one. The network manager 4 provides the client 6 with information specifying the configuration of the network. The client 6 can initiate a connection between a source device and a destination device across the network. All requests for initiating such changes to the configuration of the network are sent to the network manager 4. Control data which does not change the configuration of the network may be issued by the client 6 directly to an ENIC and, via ENIC, a group associated with the ENIC. For example control data causing the video, switcher D8 to switch may be sent directly from the client 6 to the switcher D8 via ENIC NI8.

The switching and routing client 6 which may be a computer, for example a PC, is linked to the network via a standard network interface card. The client 6 in this example controls the ENIC NI8 and/or the video switch D8 and controls the supply of video to the ENIC NI8. It also controls supply of video and audio to other destinations, e.g. D2 and D3 via ENICs NI9 and NI10. The control of the ENIC NI8 by the switching and routing client 6

takes place via network manager 4. The further switching and routing client 61, if provided, would control another destination device. The following description assumes that the network comprises only one switching and routing client 6.

Protocols and Data flows, Figure 2

The network is an Ethernet network on which various conventional protocols including UDP/IP, TCP/IP, and IGMP (Internet Group Management Protocol) are run. Other protocols are also used as will be described below.

Referring to Figure 2, Figure 2 is a simplified diagram of the network of Figure 1 showing only the Network manager 4, the client 6, and some of the ENICs: NI1, NI2 and NI8 by way of example.

AVSCP

The network manager 4 communicates with the ENICs using a protocol AVSCP (Audio Video Switching Control Protocol),which is proprietary to Sony Corporation. AVSCP uses UDP (*User Datagram Protocol*) to carry its messages. UDP is a connectionless transport protocol, which means that a one-way data packet is sent by the sending device without notifying the receiving device that the data is en route. On receipt of each data packet, the receiving device does not return status information to the sending device. The data format is described in the section Data Format and Figure 3B below.

AVSCP is a proprietary protocol used to determine the routing of audio, video and control data. AVSCP is used for communication between the network manager and each ENIC for the purpose of connection control and in order to monitor the operation status of ENIC and AV ports. For example if it is desired to connect a video tape recorder (VTR) destination to a camera source to receive AV data then the switching control server 6 (see Figure 1) must send an instruction to the ENIC port associated with the destination device, in this case the VTR, to join the specific multicast group that is sourced from the camera. This instruction between the ENIC and the switching control server 6 is sent via the AVSCP protocol.

AVSCP has five main functions which are to:

1) Monitor the operational status of the ENICs;

2) Discover the configuration of an ENIC;

3) Stop and start audio and video source transmission;

4) Direct ENICs and their associated audio and video devices to join multicast groups; and
5) Set up and tear down paths for conveying control data across the network.

The network manager 4 should be aware of the operational status of an ENIC before it can send any instructions to it. Accordingly the AVSCP requires an ENIC to send status messages to the server periodically. The network manager 4 can only control AV stream transmission and reception of an ENIC when it is operational. The network manager 4 can alternatively obtain the current configuration of an ENIC by sending a configuration request message to it. The ENIC responds by returning a message specifying the current configuration.

Figure 14 schematically illustrates how AVSCP relates to other functional modules in an ENIC network protocol stack. The arrangement of Figure 14 shows identical protocol stacks of two different ENICs 1100A and 1100B and a protocol stack of the network manager 1120. The ENIC protocol stack comprises an AVSCP layer 1104 that sits on top of a UDP/IP/Ethernet layer 1102. Other protocols 1106 may also be implemented at the same level of the protocol stack as AVSCP. The AVSCP layer communicates with a higher layer comprising ENIC applications via an AVSC_request command and an AVSC_indicate command. An uppermost layer of the ENIC protocol stack 1100A represents a local configuration 1110 of the network. The network manager protocol stack 1120 is similar to the ENIC protocol stack 1100A in that it comprises an AVSCP layer 1124 that sits on top of a UDP/IP/Ethernet layer 1122. However, a server applications layer 1128 sits on top of the AVSCP layer 1124 and communications between these two layers are mediated by the AVSC_request command and the AVSC_indicate command. The server applications layer 1128 is in communication with a higher layer corresponding to a network configuration database 1130. The AVSCP protocol layer of the ENIC 1104 may send AVSCP protocol messages to the corresponding AVSCP protocol layer 1124 of the network manager.

The AVSCP_request is a primitive command that is sent from the application layer 1108, 1128 to the AVSCP protocol layer 1104, 1124. An application initiates an AVSCP_request in order to send an AVSCP message to another AVSCP entity. The AVSCP request has the following parameters: IP address of the message destination (typically an ENIC); AVSCP message type (e.g. STOP_TX, SWITCH etc.); and number of information elements required by the message.

One or more remote client controller devices (not shown) may access the server

applications layer of the network manager 1120 via a client controller interface (not shown). The client controller interface of the network manager 1120 enables a client controller to connect remotely with and exercise a subset of control functions over a subset of ENIC devices.

Figure 15 schematically illustrates the structure of an AVSCP header. The AVSCP header has a fixed length of 32 bits. The first octet (bits 0 to 7) is used as a protocol identifier. It has a value of 0xCC. The purpose of a protocol ID is to detect possible collision with other protocols if they happen to use the same port number. The second octet (bits 8 to 15) is used as version number for the protocol. The third octet (bits 16 to 23) is reserved for future use. The fourth octet (bits 24 to 31) indicates the message type. The last 4 octets of the AVSCP header is a session-ID, which is a random number chosen by the command message initiator to tie up acknowledgement message returned by the responder to the original command message.

<u>CNMCP</u>

The network manager 4 and the switching and routing client 6 communicate with each other using CNMCP (Client Network Manager Communication Protocol) the messages of which are carried by TCP (See Section Data Format and Figure 3B below for a description of the data format). CNMCP is also a protocol that is proprietary to Sony Corporation. TCP is a connection-oriented protocol, which means that before any data is transferred between network nodes, the sending and receiving devices must co-operate in the establishment of a bi-directional communication channel. Subsequently, each package of data sent across the local network receives an acknowledgement and the sending device records status information to ensure that each data package is received without errors.

CNMCP is a protocol that is used for communication between the network manager and the clients. CNMCP enables control messages such as registration request, a switch request or a permissions update from client to network manager and further enables control messages such as a registration response, a switch response, an update indication (specifying device configurations) and a permission response from network manager to client. By sending CNMCP messages to the client, the network manager 4 informs the client 6 of data associated with the ENICs which are connected to the network and data associated with source devices and destination devices connected to the network by the ENICs. By sending

CNMCP messages to the client, the network manager 4 informs the client of the multicast IP addresses on which it can receive the proxy video streams, audio streams and status streams. The network manager can determine whether sufficient bandwidth is available to service a client request and mediates access to network resources accordingly. However, it is also possible for the client to join a multicast group directly without requesting access via network manager. This may be appropriate where, for example, the client requires only a low data-rate connection.

Alternatively to CNMCP, a known protocol such as Simple Network Management Protocol (SNMP) may be used. The client 6 can cause the network to connect Audio and Video streams from source devices specified by the client to destination devices specified by the client and to specify control data routing by sending CNMCP or SNMP messages to the network manager.

Audio and Video Data (RTP)

For sending streams of audio and video data from the sources to the destinations, the transport layer is UDP multicast. The audio and video data are carried in Real-Time Protocol (RTP) format within a UDP packet. This applies to the audio data, the full resolution video and the low resolution proxy video. (See Section Data Format and Figure 3A below for a description of the data format). RTP provides functions to support real-time traffic, that is, traffic that requires time-sensitive reproduction at the destination application. The services provided by RTP include payload type identification (e.g. video traffic), sequence numbering, time-stamping and delivery monitoring. RTP supports data transfer to multiple destinations via multicast distribution if provided by the underlying network. The RTP sequence numbers allow the receiver to reconstruct the original packet sequence. The sequence numbers may also be used to determine the proper location of a packet. RTP does not provide any mechanism to ensure timely delivery, nor does it provide other Quality of Service guarantees.

When an ENIC receives an AVSCP switch request from the network manager 4, the ENIC sends an IGMP join message to the switch 2 to join the multicast group of the data it needs to receive.

Unicast Control Data (UCD)

Control data may be sent, only as a unicast, directly from one ENIC to another. That allows, for example, a controller connected to one ENIC to control a device connected to

another ENIC. For example, commands such as play, pause stop, record, jog etc. may be sent from a controller across the network to a group such as a VTR. Such control data may also be sent from the client 6 and/or the network manager 4 to a control a device. The control channels are set up using AVSCP- see ANNEX 1. The control data itself is carried in UDP in this embodiment. However, TCP may alternatively be used to carry control data.

Stream Status (SS)

A controller connected to the network at one ENIC and controlling a group connected to a second ENIC needs to know the status of the controlled group. Status data may be sent from the controlled group to the controller via network. Also the client 6 may wish to receive the status data. Since status data is likely to be low-bandwidth, CNMCP is used to enable the client to receive the status information without the intervention of the network manager.

Data Formats- Figures 3A, 3B, 3C.

Audio and Video Data

Referring to Figure 3A, the audio and video data format comprises, in order, an Ethernet header, an IP multicast header, a UDP header, an RTP header, a field specifying the type of payload, the payload, and a CRC field. The Ethernet header comprises a source Ethernet address and a destination multicast Ethernet address. The IP multicast header comprises the source ENIC IP address and the destination multicast IP address. There are several different IP address classes e.g. Class A has the first 8-bits allocated to the network ID and the remaining 24-bits to the host ID whereas Class B has the first 16 bits allocated to the network ID and the remaining 16-bits to the host ID. Class D IP addresses are used for multicasting. The four leftmost bits of a Class D network address always start with the binary pattern 1110, corresponding to decimal numbers 224 through 239, and the remaining 28 bits are allocated to a multicast group ID. The Internet Group Management Protocol (IGMP) is an Internet layer protocol used in conjunction with multicasting and Class D IP addresses.

The set of hosts (i.e. source and/or destination devices) listening to a particular IP multicast address is called a host group. A host group may span multiple networks and membership of a host group is dynamic. The Class D IP address is mapped to the Ethernet address such that the low-order 23 bits (of 28) of the multicast group ID are copied to the low-order 23 bits of the Ethernet address. Accordingly 5 bits of the multicast group ID are not used to form the Ethernet address. As a consequence the mapping between the IP

multicast address and the Ethernet address is non-unique i.e. 32 different multicast group IDs map to the same Ethernet address.

The UDP header comprises source and destination port numbers, which are typically associated with a particular application on a destination device. Note that UDP is redundant in the case of multicast messages since in this case the multicast group address identifies the stream/content. The audio/video streams are transported using RTP protocol. Forward Error Correction (FEC) may be used for certain data streams e.g. full resolution video streams to provide a level of protection against data corruption due to network errors. FEC is provided using a known RTP payload format that provides for FEC. FEC is a parity-based error protection scheme. A known extension to the RTP protocol allows a video scan line number to be specified in the RTP payload header. The RTP header also comprises a field to specify whether 8-bit or 10-bit video is present. Although known RTP and RTP/FEC protocol formats provide the data packet fields necessary to transport audio and video data over an IP network it may also be desired to transmit additional information such as source status and source timecode information. For example if the source is a VTR then the timecode as stored on the tape should be transferred across the network. The source status information might indicate, for example, whether the VTR is currently playing, stopped or in jog/shuttle mode. This status information allows a user to operate the VTR from a remote network location. Since the timecode data and source status information is required only once per field, the information is transported in an RTP packet marked as vertical blanking. To allow audio and video resynchronisation, the RTP timecode is based on a 27MHz clock. The payload type field contains data indicating the type of payload. i.e. video or audio. The payload field contains the video or audio data. The CRC is a cyclic redundancy check known in the art.

<u>AVSCP and CNMCP</u>

AVSCP and CNMCP messages are carried by a data format as shown in Figure 3B. The format comprises in order, an Ethernet header, an IP header which is not a multicast header), a UDP or TCP header, the payload, and a CRC field. The Ethernet header comprises source and destination Ethernet addresses. The IP header comprises the source ENIC IP address and the destination ENIC IP address. UDP is used for AVSCP and TCP is used for CNMCP. The payload field contains the AVSCP or CNMCP message. The CRC is a cyclic redundancy check known in the art.

Stream Status Format

The stream status format is identical to the audio and video data format as shown in Figure 3A, with the exception of the content of the payload section. The frame comprises an Ethernet header, an IP multicast header, a UDP header an RTP header, a payload type identifier a stream status data payload and a CRC field.

Unicast Control Data Format

The unicast control data format is shown in Figure 3D and comprises an Ethernet header, a standard IP header (not multicast), a UDP header, a payload section assigned to unicast control data and a CRC field. IGMP is a known protocol. Multicasting that extends beyond a single network is complicated by the fact that Internet routers must establish whether any hosts on a given physical network belong to a given multicast group. IGMP is typically used to establish this information. IGMP lets all nodes of a physical network know the current association of hosts to multicast groups. IGMP messages are transmitted in IP datagrams and have a fixed 8-byte IGMP message size concatenated with a 20 byte IP header. The IGMP message includes a 32-bit Class D IP address.

A number of IGMP queries and reports are used by multicast routers (e.g. the switch 2 of Figure 1) to record which network interfaces have at least one host (device or group) associated with a multicast group. When the router receives a multicast message to forward, it forwards the message only to interfaces that currently have hosts associated with that multicast group.

ENIC, Figure 4

An ENIC joins a multicast group by sending an IGMP join message to the asynchronous switch 2. An ENIC may send and/or receive data in all of the formats shown in Figures 3A, 3B and 3C. It will be noted that an ENIC does not send or receive CNMCP data.

Referring to Figure 4, an ENIC comprises a network processor 20, a buffer and packet switch 22, a packetiser/depacketiser 24, a control processor CPU 26, a peripheral component interconnect (i.e. computer interface) PCI 28, a clock 202, a clock synchronisation circuit 204 and a frame synchronisation circuit 205. The clock synchronisation circuit 204 is described in co-pending UK patent Application 0204242.2. The frame synchronisation circuit is described in co-pending patent application ____________ (Agent reference P015702GB)

The packetiser/depacketiser has 3 video inputs 218 for receiving respective SDI video streams, 3 audio inputs 220 for receiving respective audio streams. Alternatively, three input ports could be provided for receiving combined SDI audio/video streams and the audio and video streams could be subsequently separated to form 3 audio and 3 video streams with in the ENIC The packetiser/depacketiser 24 has likewise 3 video outputs 222 and 3 audio outputs 224.

The CPU 26 has 3 control data inputs 226 and 3 control data outputs 228.

The three video inputs 218 are connected to respective ones of three substantially real-time proxy video generators 212 which generate the low resolution versions of the video streams as will be described below. The outputs of the proxy generators and the inputs 218 are connected to a packetiser and multiplexer 214, which converts the full resolution serial video from the inputs 218 and the proxy video from the proxy generators 212 to packets. The packets are supplied to the packet switch 22. The packetiser/depacketiser 24 has a depacketiser 216 and demultiplexer for receiving packets representing the video and audio channels from the packet switch 22. It depacketises and demultiplexes the video and audio into 3 serial video streams and 3 serial audio streams for supply to respective ones of 3 video outputs 222 and 3 audio outputs 224. Thus the packetiser/depacketiser 24 provides routing of the video and audio received from the packet switch 22 to outputs 222 and 224 and routing of the video and audio received from the inputs 218 220 to the switch 22. The packetiser/depacketiser 24 also provides synchronisation of the different video and audio streams in conjunction with the clock synchronisation circuit 204 and frame alignment of the video frames of the different video streams in conjunction with the frame synchronisation circuit 205.

The packet switch 22 provides routing of video, audio and control packets received from the network processor 20 in accordance with tags applied to the packets in the network processor 20. The network processor 20 generates the tags in accordance with header data in the received packets. There are two sorts of tag: a "flow" tag which defines the route through the packet switch 22 and a "type" tag which defines the final output to which the packets are supplied by the packetiser/depacketiser 24. The video and audio packets are routed to the depacketiser 216 and the control packets are routed to the CPU 26. The CPU provides 3 control data channels 228 here denoted "*RS422*" because they provide control similar to that

provided by RS422 in a conventional studio. CPU 26 also receives 3 control data channels 226.

The network processor 20 comprises UDP/IP filters 208, which detect sync, audio, video, status and control data packets received from the network using the headers of those packets. Clock sync packets are directed by the network processor 20 directly to the clock synchronisation circuit 204 to synchronise the clock 202 with a master reference clock as described in the co-pending UK patent application 0204242.2. Frame sync packets are directed by the network processor to the frame sync circuit 205. The network processor directs the sync packets directly to the clock synchronisation circuit 204 frame synchronisation circuit 205 to reduce time delays which would reduce the accuracy of the synchronisation. Other packets, for example AVSCP, which are not recognised by the filters 208 are directed to the CPU 26 although filters could be set up for these.

Tags are attached to the audio and video packets in accordance with the header data received with them. The tagged video and audio packets are supplied to the packet switch 22 which routes them to the depacketiser 216 or PCI 28. The tagged control data packets are directed by the packet switch to the CPU 26. The packet switch 22 is described in more detail below.

Routing Data in an ENIC

1. Data received from the network

An ENIC may receive from the network 2: audio and video data as shown in Figure 3A; AVSCP data as shown in Figure 3B; stream status data (which is in essentially the same format as shown in Figure 3A); and unicast control data as shown in Figure 3D. The Ethernet header provides the physical address of the ENIC allowing a packet to be delivered by the network 2 in known manner to the ENIC.

The network processor 20 of the ENIC (see Figure 4) has the UDP/IP filters 208 that extract the IP and UDP headers, decode the address information in the headers, detect the payload data type from the payload type field (see Figure 3A). The network processor 20 then replaces the packet header with a tag identifier, which specifies a data processing route through the ENIC for the corresponding data to a target data handling node such as a video or audio processor. Figure 5A schematically illustrates the data format of a tagged packet. The tagged data packet is 32 bits wide and is of indefinite length i.e. the payload has a variable length. The first 32 bits of the tagged packet comprise an 8 bits "flow" data field, an 8-bit

"type " data field and a 16-bit "size" field. The next 32 bits are currently unused. The unused filed is followed by a payload field. For audio and video data the tagged packet payload comprises the RTP header and the payload type data in addition to the audio or video data payload of Figure 3A. In the case of AVSCP/CNMCP data packets and unicast control data packets (see Figures 3B and 3C) the tagged packet payload is the message data.

The flow data field defines the output of the packet switch 22 (Figure 4) corresponding to the target data-handling node for which the tagged packet payload is destined. The type data field determines what the target processor does with the data and the size data field specifies the payload size.

Figure 5B schematically illustrates an example of a flow assignment allocation. In this example flow 0 corresponds to data that will not be passed to any target processing device e.g. untagged data; flows 1 and 4 correspond to video input and output ports 218, 222 (see Figure 4); flows 2 and 5 correspond to CPU data flows from and to the network; and flows 3 and 6 correspond to PCI data flow from and to the network.

Figure 5C schematically illustrates how video data, PCI data, network data and CPU data is mapped to each of the six defined flow paths via multiplexers (MUX) and demultiplexers (DEMUX). Each of the data flows of Figure 5B is associated with a FIFO. There is no direct means of determining the size or number of packets written to the FIFO since this is not necessary. The tags associated with the packets specify the packet size so only a "not empty" indication for the FIFO is required by a MUX to perform a read operation. The MUX modules are programmable (by external means such as a CPU) such that they are sensitive only to particular flows. This enables virtual flow paths to be set up across the buffer and packet switch 22 of Figure 4. Similarly, to avoid contention, only a single DEMUX module can write into any one data flow. Again, the mapping is programmably controlled by external means.

Referring to Figure 6A, the video section of the packetiser/depacketiser 24 is shown. It comprises a demultiplexer 2401 which responds to the "type" data in the tags attached to the video packets to feed video packets to three channels V0, V1 and V2 denoted by the type data. Each channel comprises an RTP/FEC decoder 2402, 2403, 2404 followed by a frame store 2405. The RTP decoder 2402 removes the tag from the packet it receives and writes the packet into the frame store at the address defined by the RTP packet header, in

particular the line number data thereof, to create a video frame having the video data in the correct order.

Example of operation 1

In this example, two video streams are sent to an ENIC for example NI8 in Figure 1 and provided by that ENIC to the video switcher D8 which switches from one stream to the other.

The ENIC receives from the network 2, video packets from two video source devices for example video outputs of camera 1 S1 and camera 2 S2. To receive those packets the ENIC NI8 is ordered by an AVSCP message from the network controller 4 to join the multicast groups of the two video source devices of groups S1 and S2. The AVSCP message is detected by the network processor 20 and fed to the CPU 26 which issues a IGMP join message to the network. The video packets are received by the network processor 20. The destination IP address for the video packets, specified by the IP header corresponds to the multicast group that defines the data stream and thus determines to which subset of destinations on the network the data stream is routed. The headers are removed by the network processor and replaced by the tags. The packet switch 22 routes the video packets to the demultiplexer 2401 according to the flow data in the tag. The demultiplexer routes the packets to channels 2402 and 2403 (by way of example) in which the video frames are reconstructed from the packets in frame stores 2405 and 2406. In addition, the frame sync circuit 205 of Figure 1 aligns the frames of the two video streams. The video switch D8 receives the two video streams from the ENIC. The video switch also receives control data from the CPU 26 of the ENIC NI8. The control data is sent by the switching and routing client 6 as unicast control data which is received by the network processor 20. The unicast control data has a header that identifies it as a control packet and these control packets are routed to a CPU on the ENIC. The control data orders the video switcher D8 to switch its output from one of the video streams to the other.

Example of operation 2.

In this example two video streams are sent to an ENIC, for example ENIC NI8 in Figure 1, and that ENIC switches an output from one stream to the other. A small network is Shown in FIGURE 7 that consists of two cameras ('Camera 1' and 'Camera 2'), a DME unit, an AB switch and a monitor with tally that shows the output of the system. The broken lines represent a network connection and the unbroken lines an SDI connection to an ENIC.

Sources are categorised as either linked sources or pure sources. A linked source is a source that has been supplied to the network by a destination device, for example a data processor (e.g. chroma keyer) that receives data via the network and subsequently outputs processed data onto the network as a source for another network device. For such linked sources, the meaning of the data (i.e. what the data conveys) is not changed by the processing operation. A pure source is a source that supplies to the network for the first time.

The video source device for each camera is a 'pure' source. Each camera has a user group called "Producer" and the Producer has set the tally tag to be the name of the cameraman, i.e. FRED or JIM. There are three other ENICs on the network. The first performs a visual effect on the video from 'Camera 1' and is called ENIC_DME since it performs digital multi-effects (DME). This ENIC will have two device entries in the database, called 'DME In' and 'DME Out'. 'DME In' is a destination device, which has a video link to 'DME Out'. 'DME Out' is a source device, which has a video link to 'DME In' and transmits onto the network. 'DME Out' also has a tally entry of El (indicating EFFECT 1). The second ENIC performs a seamless switch between two sources and is called ENIC_AB_SWITCH. This ENIC will have three device entries in the database, called 'Switch A In', 'Switch B In' and 'Switch Out'. 'Switch Out' is a source device that will either be linked to 'Switch A In' or 'Switch B In', depending on which video source is selected. The third ENIC is called ENIC_AIR and has one device called 'Monitor' (a monitor with a tally). 'Monitor is a 'pure' destination device (LINKED = 0) that receives video from the AB switch. It also has a tally that displays the METADATA from its source device, 'Switch Out'.

Say that the AB switch is showing channel A and the METADATA of Camera 1 changes. Maybe ROB replaces FRED. A request will be sent to the Network Manager to changed FRED to ROB. The Network Manager will examine each destination Device that is subscribed to Camera 1's multicast group and update the view of any client that is displaying it. If any of these destination devices is a linked device, then it must navigate to the corresponding linked source device and update all of its destinations, and so on. In the scenario above, 'DME In' is the only destination device and it is linked to 'DME Out'. 'DME Out's' METADATA (El) is concatenated to ROB to form ROB_El and all of its destinations must be notified. The only destination device is 'Switch A In'. Since the switch is showing channel A, 'Switch A In' is a linked destination device and we must update all destinations of

its corresponding source device ('Switch Out'). 'Switch Out' only has one destination and this is a pure destination 'Monitor'. The tally of 'Monitor' is updated with 'ROB_El'.Another scenario is when we perform an AB switch. A request will be sent to the Network Manager to perform a seamless AB switch between devices 'Switch A In' and 'Switch B In'. If the network has been configured correctly, as above, then this is allowed. Since SRC_DEVICE of 'Switch B In' references 'Camera 2', we can navigate to 'Camera 2' and update its status as being *ON AIR*. Similarly, we can take 'Camera 1' *OFF AIR* by navigating back through the devices from 'Switch A In'. The correct tally, i.e. JIM can now be propagated to 'Monitor' as already described.

2. Sending data to the network Figures 6B and 6C.

Referring to Figure 6B, one channel of SDI video is received by a buffer 2408 from an SDI source such as a camera. The buffer 2408 stores the video temporarily whilst it is packetised by an RTP/FEC encoder 2410 and sent to a buffer 2412 for temporary storage. A tag generator 241 adds to the RTP packets a tag comprising flow and type data as shown in Figure 5. A multiplexer 2416 receives the tagged packets from the tag generator and multiplexes the video packet with other video packets from similar video channels. The tag is defined by data produced by the CPU 26 in response to an AVSCP message received from the network controller 4. As shown schematically in Figure 5B, the packet switch directs the video packets to the network processor (net) or to the PCI 28 according to the flow data in the tag. Audio packets are similarly processed and routed.

Where packets are to be routed to the network, a header generator 210 strips the tag from the packet and, based on the flow and type flags, generates an appropriate part of the network header which is appended to the packet.

Proxy Video

Referring to Figure 8, proxy video is generated from SDI video as follows. A horizontal filter 70 applies a low-pass FIR filter to the SDI input data. Output of the horizontal filter is supplied as input to a horizontal subsampler 71, which subsamples the SDI video horizontally to reduce the horizontal resolution. A vertical subsampler 72 reduces the vertical resolution of data received from the horizontal subsampler 71. The resulting proxy video is then encoded by an encoder 74 to form RTP packets. There is one proxy video

generator for each video channel. The proxy video is processed in the same way as the SDI video by the packetiser 24, the packet switch 22 and the network processor 20. The proxy video is always directed to the switching and routing client 6, or one of the switching and routing clients 6 and 61. Thus one proxy video stream is multicast in a first multicast group to which the client 6 and/or 61 joins and the SDI video from which that proxy video is derived is multicast in a second multicast group. The multicast group is defined by the class D IP address that identifies the data stream. In an alternative embodiment alternate fields of either the proxy video stream or the higher-resolution SDI video stream could be assigned to different multicast groups.

In a currently preferred example of the invention, the proxy video comprises180 samples × 144 lines (PAL) or 180 samples ×120 lines (NTSC) and 25 or 30 frames per second, with horizontal and vertical filtering. The number of bits per sample may be 24 bits (i.e. 3 colours, each 8 bits) or 16 bits (i.e. 3 colours, each 5 bits).

<u>Switching and routing client 6</u>

Referring to Figures 9 and 10, examples of graphical user interfaces (GUI) are shown. In this example of the invention, the GUI is provided by the switching and routing client 6. However the GUI may be provided by the network manager 4 or by both the network manager 4 and the switching and routing client 6. The GUI is an interface with underlying software which reacts to actions taken by the user using the GUI.

<u>Data flows</u>

The GUI displays information about the configuration of the network provided to it by the network manager 4. That information is provided using the CNMCP protocol as discussed above. The GUI also displays proxy video provided by the ENICs using the Real Time Transport protocol described above. The proxy video is multicast and to receive it the switching and routing client joins the multicast groups of the proxy video streams. The routing of data is established using IGMP message commands. The GUI may be used to initiate control of a controllable group such as a VTR. The switching and routing client 6 unicasts control data directly to the ENIC associated with the controlled group in response to an action taken via GUI. Unicast control data is described above. The switching and routing client receives status stream data which is multicast as described above and the client joins the multicast group thereof.

When the GUI is used to initiate a routing of video from a source device to a destination device, it sends an CNMCP message to the network manager 4. The network manager sends an AVSCP message to the ENIC associated with the destination device to cause it to join the destination device to the required multicast group.

The client 6 is able to send IGMP join messages to the network. However, the client may also self-subscribe to a multicast group for communication of status, audio and proxy data streams only. The network manager controls client access to a multicast group corresponding to a video stream. The GUI

The following description assumes that the GUI is operated in conventional manner using at least a pointing device such as a mouse and/or a keyboard. Alternatively, a keyboard interface having "hot keys" mapped to particular GUI commands or a touchscreen interface may be used to issue commands. The GUI of Figure 9 has three main display areas A1, A2, and A3.

Area A1 is a network management area displaying graphical representations of the groups (e.g. Cameras CAM1 etc and VTRs VTR1 etc) and their source devices (e.g. output CAM V1 of CAM1). The graphical representations of the groups are displayed with the tally text (e.g. CAM 1) and the source devices with their sub-tally text (e.g. CAM V1). The data for creating the display in the area A1 is derived from the database held by the network manager and provided to the switching and routing client using CNMCP messages.

Area A2 is a source content review area which has plurality of proxy video display areas or windows W1 to W10. In this example there are 10 such windows but there can be any convenient number thereof. The windows W1 to W10 display proxy video. In this example the proxy video to be displayed in the windows is chosen by dragging a source device from the network management area A1 to a window. That causes the underlying software to send an IGMP join message to the network to join the required multicast group

The windows have respective label areas L1 to L10 in which the GUI displays the appropriate tally text and/or sub tally text.

Area A3 is a routing review area comprising buttons B which act as switches. There are two rows of buttons in this example: a row of buttons associated with groups and/or source devices and labelled with the appropriate tally text and a row of destination buttons also labelled with tally text. By operating a source button and one or more destination buttons the source indicated by the button is linked across the network to the selected destinations. In

the example of Figure 9, the highlighted buttons show CAM 1 is linked to MON1, VTR 2, and DSP2. Audio associated with CAM1 is linked to AU OUT 3.

By way of further explanation, assume a source CAM1 is to be connected to MON1. When a Network Client starts up, it connects to the Network Manager on a known port and the Network Manager sends information on the source devices, which are available to it. This allows the client to form a view of the network. Each device is delivered to the client with an ID, which the client will use to describe the device in subsequent communication with the Network Manager. A destination device may be a Monitor for example. If the client wishes to route video from a source group (e.g. a VTR) then it will send to the Network Manager a CNMCP Switch message that contains the IDs of the destination device and of the source device.

If the client is not permitted to perform this operation then the Network Manager will send a CNMCP NAK message to the client in response. Otherwise it will process the request as follows.

The Network Manager will examine the database and determine which multicast IP address the video is being transmitted to. An AVSCP Switch message will be created that contains this IP address and it is sent to the ENIC, which connects the Monitor. The ENIC embedded software sends an IGMP Join message to this IP address and sends an AVSCP ACK message back to the Network Manager. The ENIC should now be receiving the desired video data and will send it to the SDI output that connects the Monitor. Meanwhile, the Network Manager, having received the AVSCP ACK message, will update the routing information in the database. The Network Manager sends a CNMCP ACK message back to the client to indicate success.

The GUI of Figure 9 preferably also includes, as shown, two further display areas M1 and M2 showing the video displayed on play-out monitors MON 1 and MON2. In this example MON2 has a dark border indicating that it shows the video being played out on LINE OUT 1 from for example VTR1. MON 1 which has a lighter border shows the video from CAM1 which has been preselected for play-out next. The video may be selected for display in the windows MON1 and MON2 by dragging and dropping proxy video from windows W1 to W10 into MON 1 and MON 2. The video to be played out may be selected by clicking on MON1 or MON2.

The GUI of Figure 9 has an audio control display area AUD.

The GUI also has virtual controls C1 to C10 associated with the windows W 1 to W10 and controls CM associated with the MON1 and 2 windows. Operation of such a control causes the underlying software to send unicast control data UCD across the network directly to the source group from which the video in the associated window originates. Alternatively, or in addition, C1 to C10 can indicate the current status of the relevant device.

The GUI of Figure 10 differs in minor ways from that of Figure 9. It has proxy video display areas W1 to W8, a network management area A1 (shown only schematically) identical to that of Figure 9, and monitor displays M1 and M2. The GUI of Figure 10 lacks the rows of source and destination buttons, but has two buttons M1 and M2 which act, like the buttons of Figure 9, as switches. The buttons M1 and M2 select for play-out video associated with an associated one of windows M1 and M2. The played out video is displayed in a play-out window MP.

The windows M1 and M2 have associated audio controls A1 and A2 for switching on and off an audio monitor to allow the user to monitor the audio associated with the video of windows M1 and M2.

The video to be displayed in windows M1 an dM2 is dragged and dropped into those windows from the proxy video windows W1 to W8. That causes the full resolution video to be sent from the sources to full resolution monitors such as MON1 and MON2 in Figure 1 and to a video switcher such as D8 in Figure 1 via ENIC NI8. Operating the button M1 or M2 causes the underlying software to send unicast control data UCD via the ENIC NI8 to the video switcher causing the video switcher to switch.

Figure 11 schematically illustrates a GUI that presents the operator with an overview of the network configuration. The GUI comprises a first source panel 110 that displays active sources and inactive sources belonging to the IP network. Source groups such as cameras CAM1, CAM2, CAM 3 are represented. The video tape recorder group VTR1 has separate audio VTR A 1/ 2 VTR A 3/ 4 and video VTR V1 devices associated with it, which are also displayed. Both the source type e.g. MIC1 for a first microphone and the source name MIC A1/2 that specifies the audio channel device are represented in the first source panel 110. The source type is represented by an icon but the source name is not. An input may be selected by highlighting a desired source on the first source panel 110, for example camera 1 CAM 1 is currently selected. A network review panel 112 comprises three sub-panels: a controllers sub-panel 114, a source sub-panel 116 and a destination sub-panel. The connection between

a controller, a source and one or more destinations is represented by colour-coded branch connections between entries in the three sub-panels. The current configuration shows that a first controller CONT 1 is controlling the source group CAM1, which in turn is providing data to six different destination devices i.e. two monitors MON1, MON2, VTR1, an audio output AUDIO OUT 3, a digital signal processor DSP 2 and an output line LINE OUT 1.

Figure 12 is one way of indicating the connections between sources and destinations across the network. Area 120 depicts groups (e.g. CAM1) and the associated source devices (e.g. V1, V2) and area 122 denotes destinations. Each source group has a coloured bar 124 associated with it. Area 121 is a matrix which uses the coloured bars to indicate the connections between sources and destinations. The source sub-panel 116 provides pull-down menus for each source that provide more detailed information about devices e.g. audio and video data streams associated with that source. The relationship between a source and digital signal processors (DSP) is indicated by colour coding in the left hand margin of the source sub-panel 116, for example in this case CAM 1 is associated with both DSP 2 and DSP 3. The names of the sources e.g. CAM 1, VTR 1, MIC 1 are derived from the tally text.

Figure 12 schematically illustrates a GUI that provides the user with an overview and an interface for displaying to an operator how the data is being routed across the network. The GUI comprises a routing review overview panel 121 at the top of the screen and a main routing review panel 122 comprising a source sub-panel 123 and a destination sub-panel 124. The overview routing review panel 121 provides an easily comprehensible overview of the relationships between sources and destinations. This is achieved by colour-coded highlighting. This panel 121 currently indicates that source CAM1 is connected to destinations MON1, MON2, MON3, VTR2 and AUOUT3. By clicking on a given source area of the routing review overview panel 121, that source and any destinations associated with it are highlighted. The sources sub-panel 124 provides an expanded view of the source in which both the source group e.g. CAM1 and the associated device V1 or V2 are graphically represented. Similarly, the destinations sub-panel provides an expanded view of the destination groups. From the highlighted areas in the sources sub-panel 121 and the destinations sub-panel 124 it is apparent that CAM1 device V1 is connected to devices V1 and V2 of MON1 for example. The destination sub panel 124 also provides a graphical colour coded matrix representation of source-destination connections.

Figure 13 schematically illustrates a user interface provided on the network manager via which a user may manually enter configuration data. When a device is connected to the network, the user informs the network manager that this is the case via the user interface. The interface comprises an ENIC ID dialog box, a PORT ID dialog box and a TALLY TEXT dialog box. The user enters into dialog boxes data required by the manager to determine the configuration of the network. The ENIC ID entry is a user-defined identifier e.g. ENIC6, the PORT ID entry specifies the ENIC port to which the device has been connected and the TALLY TEXT entry specifies the freely assignable label (referred to above as tally text) used as a source/destination identifier. The tally text ID is used in addition to (rather than as an alternative to) the source and destination identifiers ID discussed above.

References.

1, RTP Payload Format for BT.656 Video Encoding, D. Tynan, (Claddagh Films) RFC2431, Oct. 1998.

CLAIMS

1. A video processor connectable to a data network and being arranged to receive a digital video data stream of a first resolution, the processor comprising:

a video generator arranged to produce from the video data stream of the first resolution a video data stream of a second, lower, resolution;

a packetiser operable to format at least that part of the video data stream of the first resolution not represented by the video data stream of the second resolution into data packets and to format the video data stream of the second resolution into data packets; and

a network interface arranged to launch the data packets onto the data network;

the video generator, the packetiser and the network interface operating substantially in real time.

2. A processor according to claim 1, comprising a multiplexer for multiplexing together the data packets corresponding to the video data streams of the first and second resolutions.

3. A processor according to claim 1 or claim 2, in which the network interface is arranged to launch the data packets corresponding to the video data stream of the first resolution onto the network as a multicast group.

4. A processor according to claim 3, in which the network interface is arranged to launch the data packets corresponding to the video data stream of the second resolution onto the network as a second multicast group, different to the first multicast group.

5. A processor according to claim 3 or claim 4, in which the network interface is arranged to format the data packets into multicast IP/UDP (Internet Protocol / User Datagram Protocol) packets to launch onto the network.

6. A processor according to any one of the preceding claims, in which:

the processor is arranged to receive two or more input digital video data streams;

the video generator is arranged to produce two or more video data streams of the second resolution from respective input digital video data streams; and

the packetiser is arranged to format at least the parts of the video data streams of the first resolution into data packets and to format the video data streams of the second resolution into data packets.

7. A processor according to any one of the preceding claims, in which the or each video data stream of the second resolution comprises an uncompressed video data stream.

8. A processor according to claim 7, in which the or each video data stream of the second resolution conforms with the RGB 555 format, representing each of red, green and blue pixel values by 5 bits of a pixel word.

9. A processor according to claim 8, in which the packetiser is arranged to format the or each video data stream of the second resolution into RTP (Real time Transport Protocol) packets having a video line of data in each such packet.

10. A processor according to any one of claims 1 to 7, in which the packetiser is arranged to format the or each video data stream into RTP (Real time Transport Protocol) packets.

11. A processor according to claim 10, in which the RTP packets carrying the video data stream of the first resolution conform to the BT.656 video encoding standard.

12. A processor according to any one of the preceding claims, in which the or each video data stream of the first resolution comprises 625 lines or 525 lines per frame by 1440 samples per line.

13. A processor according to any one of the preceding claims, in which the or each video data stream of the second resolution comprises either:

- 144 lines per frame by 180 samples per line in the case of a video data stream of the first resolution having 625 lines per frame; or
- 120 lines per frame by 180 samples per line in the case of a video data stream of the first resolution having 525 lines per frame.

14. A processor according to any one of the preceding claims, in which the or each video generator comprises a first filter and subsampler for subsampling the video data in the line direction and a second subsampler for subsampling the video data in the frame direction.

15. A processor according to any one of the preceding claims, in which the packetiser is arranged to format the video data stream of the first resolution into data packets.

16. A processor according to any one of claims 1 to 14, in which:

a subset of the video data stream of the first resolution can be derived from the respective video data stream of the second resolution; and

the packetiser is arranged to format only a part, being all but the subset, of the video data stream of the first resolution into data packets.

17. A processor according to any one of the preceding claims, in which:

the processor is arranged to receive one or more audio and/or control data streams associated with the digital video data stream;

the packetiser is arranged to format the audio and/or control data streams into data packets; and

the network interface is arranged to launch the data packets of the audio and/or control data streams onto the network.

18. A processor according to any one of the preceding claims, in which:

the video data stream of the first resolution represents an interlaced video signal; and

the video generator is selectably operable to generate the video data stream of the second resolution from only odd fields or from only even fields of the video data stream of the first resolution.

19. A processor according to any one of claims 1 to 17, in which:

the video data stream of the first resolution represents an interlaced video signal; and

the video generator is operable to generate a video data stream of the second resolution from only odd fields and a video data stream of the second resolution from only even fields of the video data stream of the first resolution.

20. A processor substantially as hereinbefore described with reference to the accompanying drawings.

21. Video source equipment comprising a video processor according to any one of the preceding claims.

22. Video destination equipment comprising a video processor according to any one of claims 1 to 20.

ABSTRACT

VIDEO PROCESSING

A video processor connectable to a data network and being arranged to receive a digital video data stream of a first resolution comprises a video generator arranged to produce from the video data stream of the first resolution a video data stream of a second, lower, resolution; a packetiser operable to format at least that part of the video data stream of the first resolution not represented by the video data stream of the second resolution into data packets and to format the video data stream of the second resolution into data packets; and a network interface arranged to launch the data packets onto the data network, the video generator, the packetiser and the network interface operating substantially in real time.

(Fig. 2)

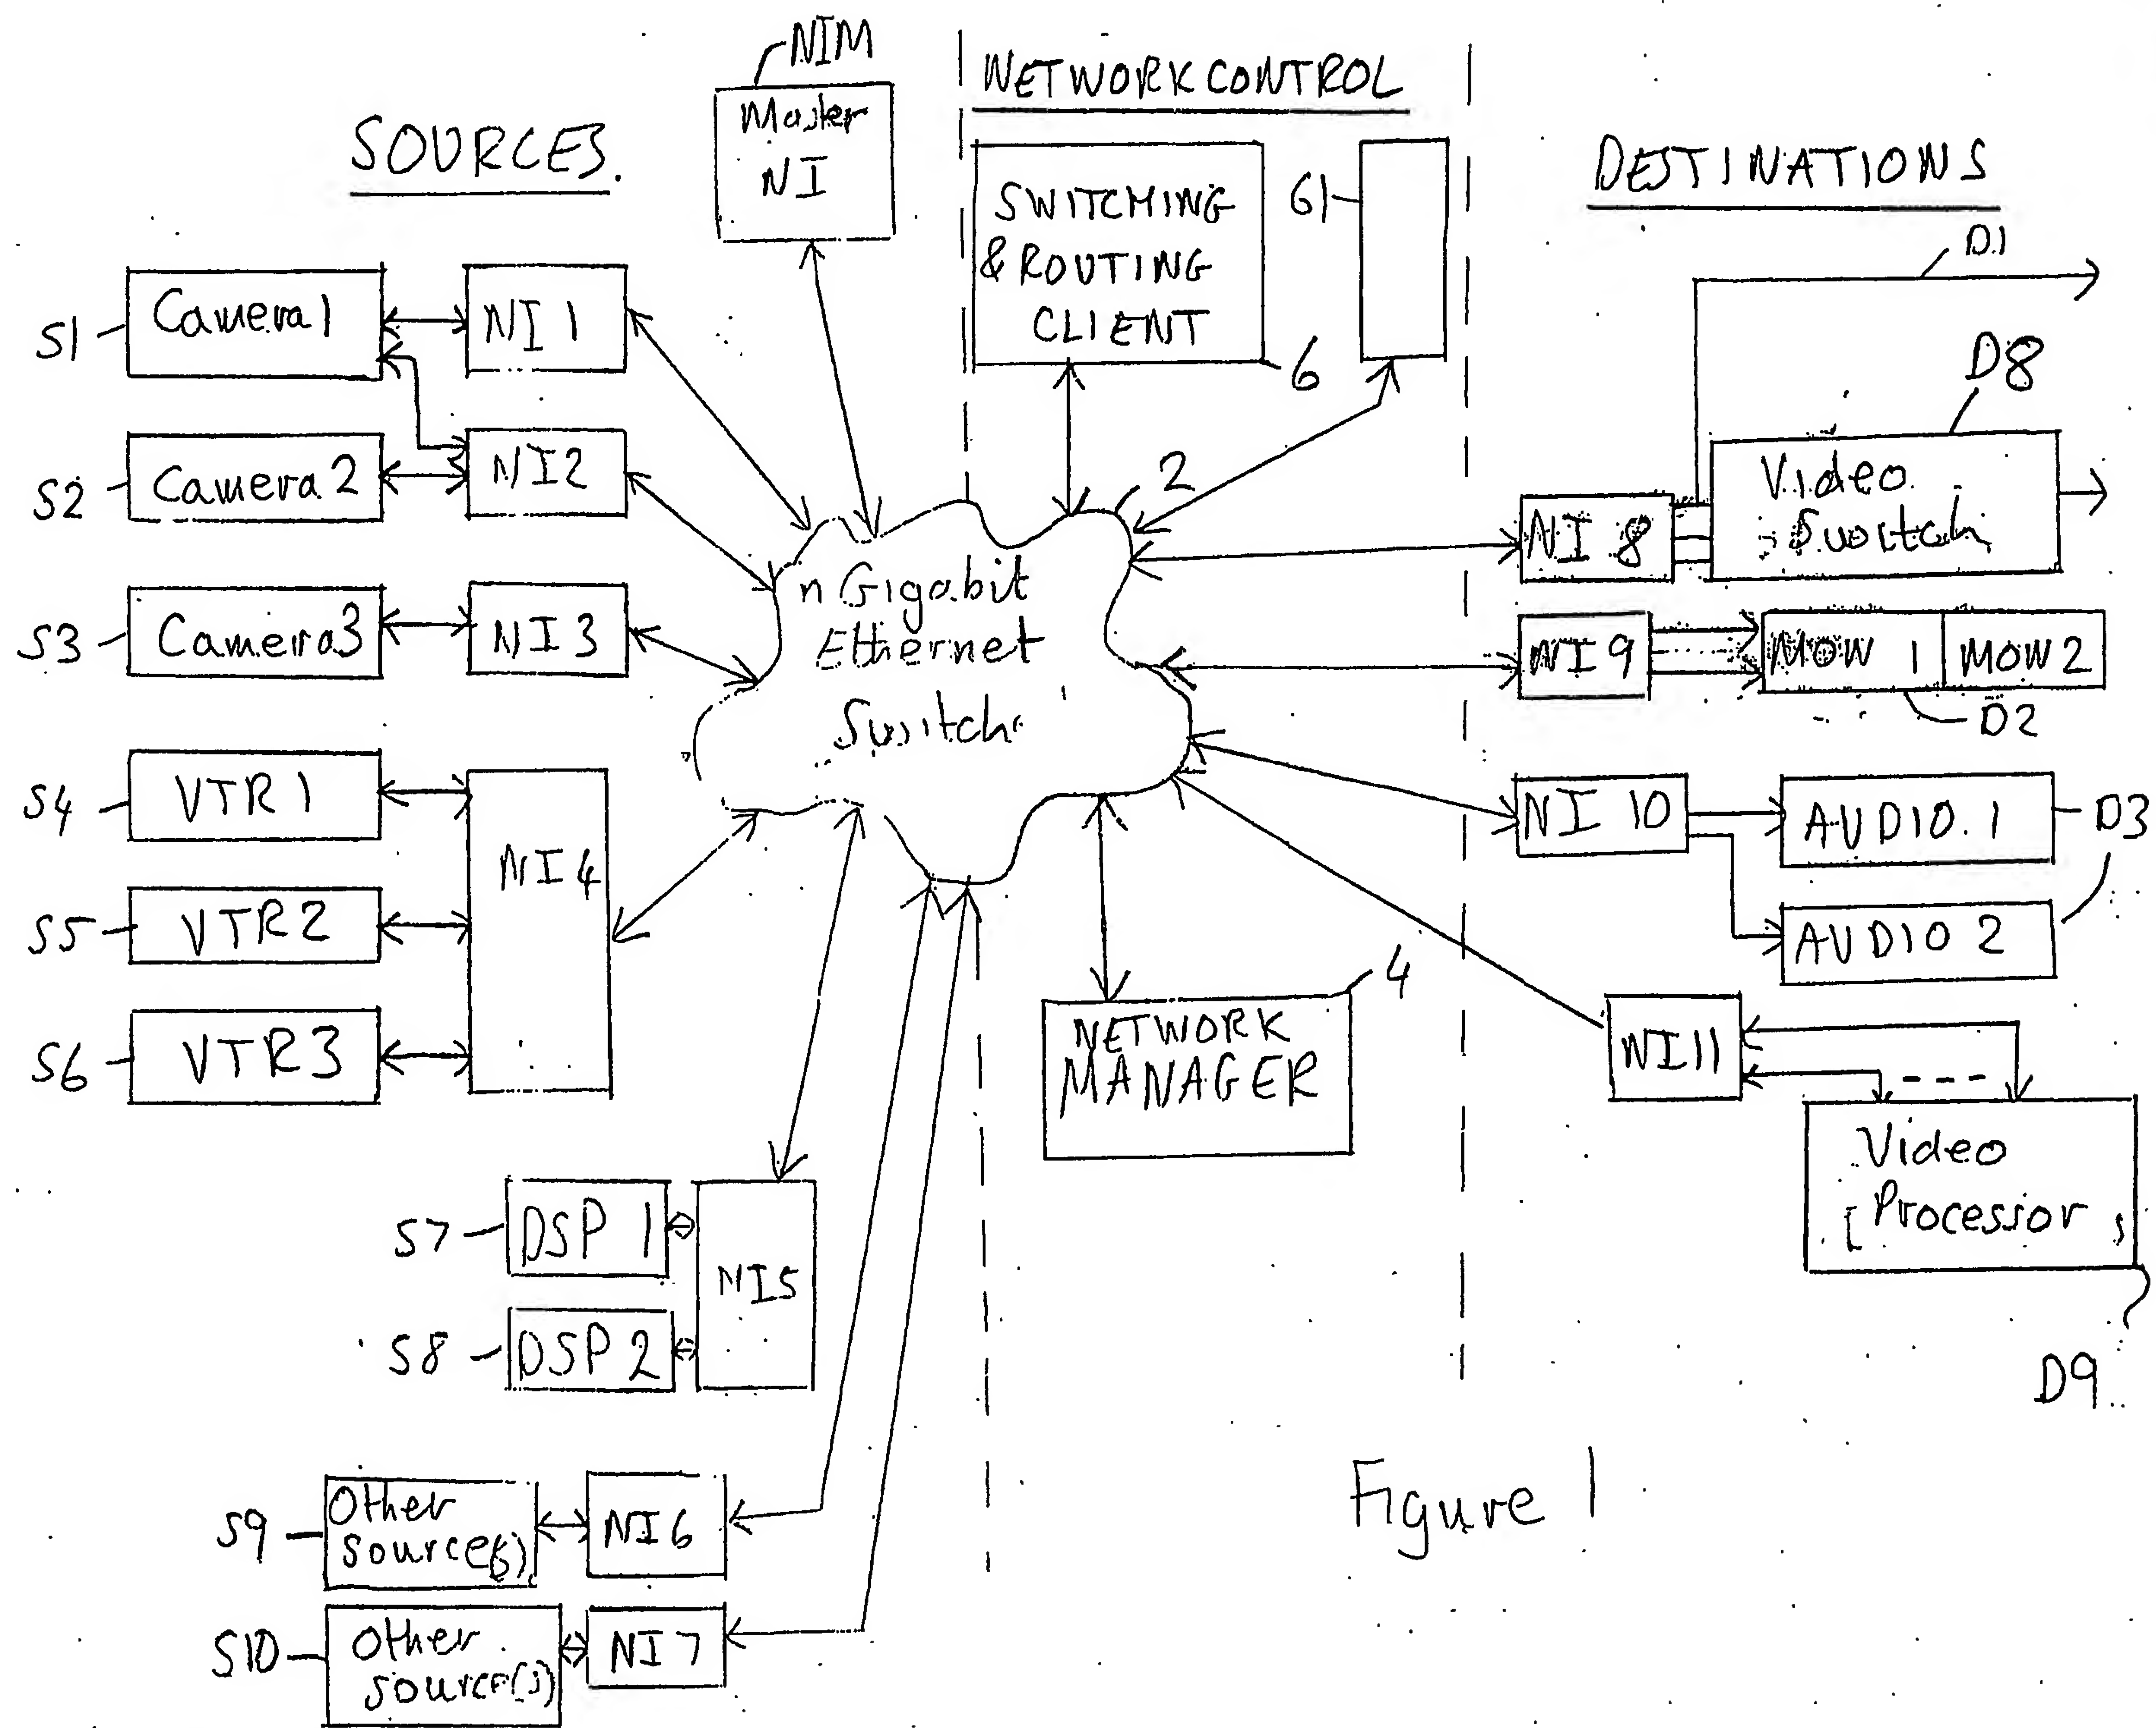

Figure 1

I-02-134.

2/10

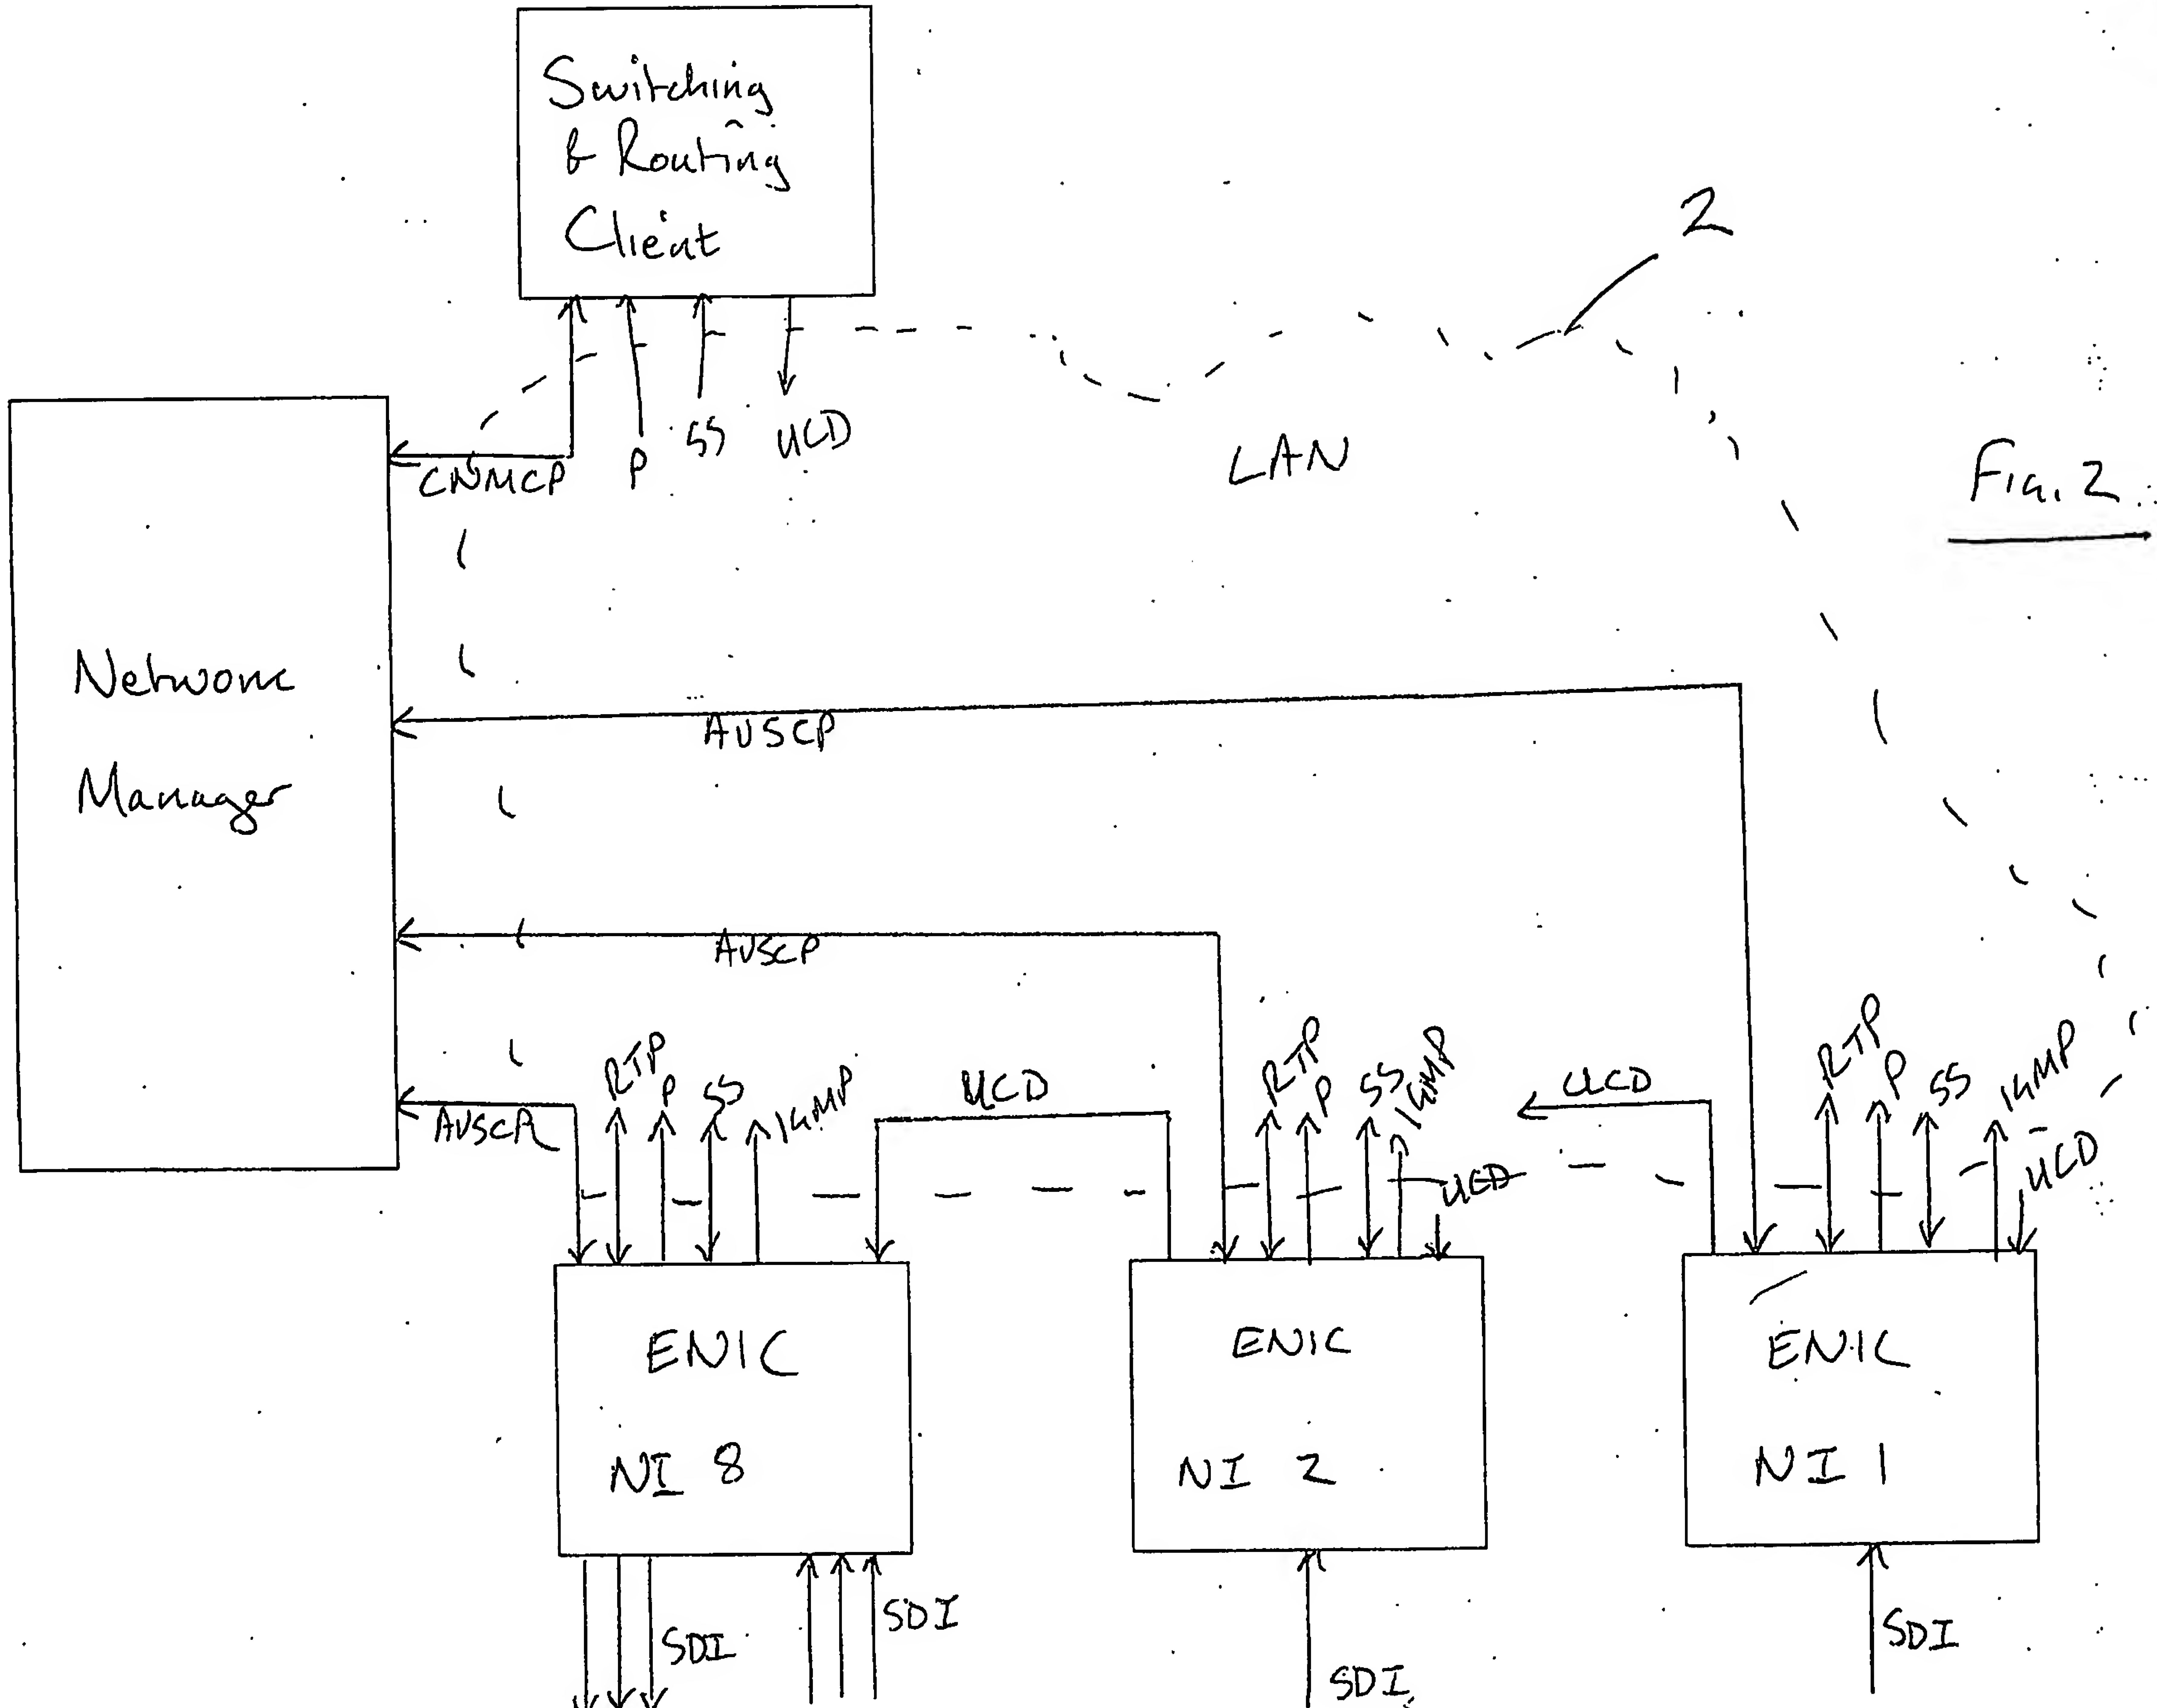

Switching & Routing Client
CNMCP
P
SS
UCD
LAN
2
Network Manager
AVSCP
AVSCP
AVSCP
RTP
P
SS
IGMP
UCD
ENIC
NI 8
ENIC
NI 2
ENIC
NI 1
SDI
SDI
SDI
SDI

Fig. 2

ETHERNET HEADER	IP MULTICAST HEADER	UDP HEADER	RTP	PAYLOAD TYPE	AUDIO/VIDEO PAYLOAD DATA	CRC

FIG 3A AUDIO/VIDEO

ETHERNET HEADER	IP HEADER (NOT MULTI-CAST)	UDP/ TCP HEADER	MESSAGE	CRC

FIG 3B AVSCP/CNMCP

ETHERNET HEADER	IP HEADER (NOT MULTICAST)	UDP	MESSAGE	CRC

FIG 3C

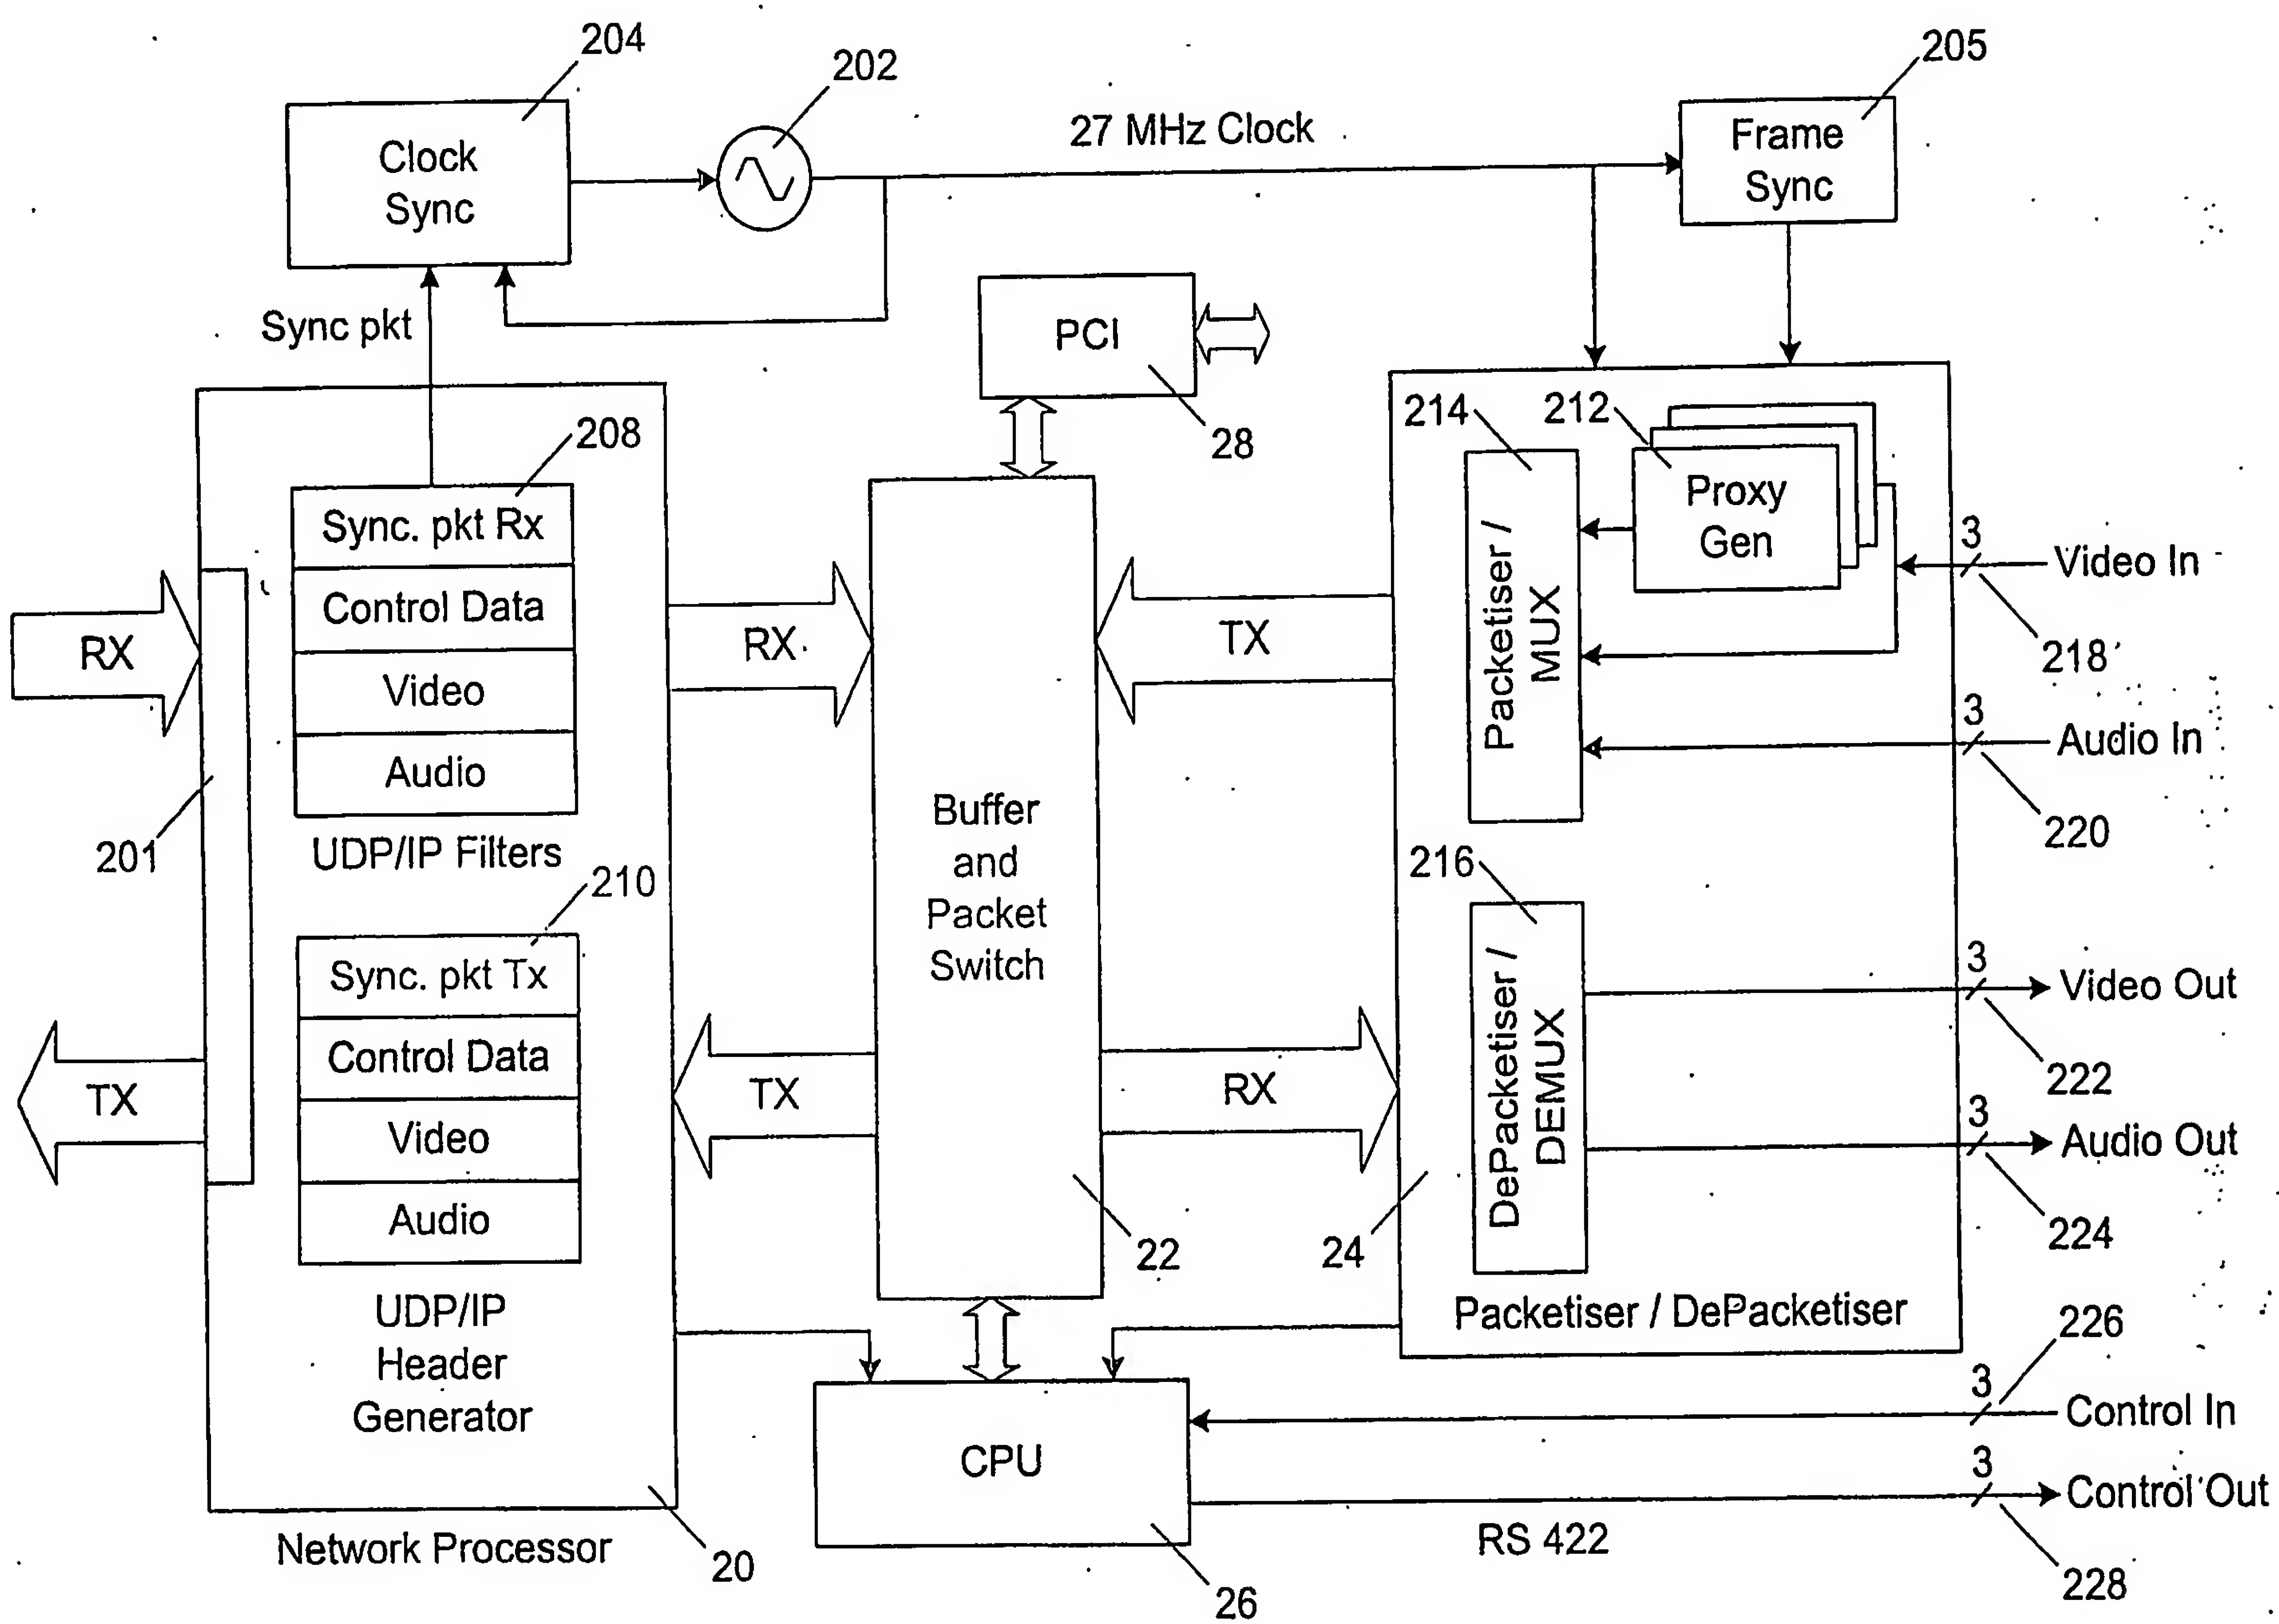

Clock Sync
204
202
27 MHz Clock
Frame Sync
205
Sync pkt
PCI
28
Sync. pkt Rx
Control Data
Video
Audio
208
UDP/IP Filters
RX
TX
201
Buffer and Packet Switch
22
Packetiser / MUX
214
Proxy Gen
212
Video In
218
Audio In
220
DePacketiser / DEMUX
216
Video Out
222
Audio Out
224
Packetiser / DePacketiser
24
Sync. pkt Tx
Control Data
Video
Audio
210
UDP/IP Header Generator
Network Processor
20
CPU
26
Control In
226
Control Out
228
RS 422
3

FIGURE 4

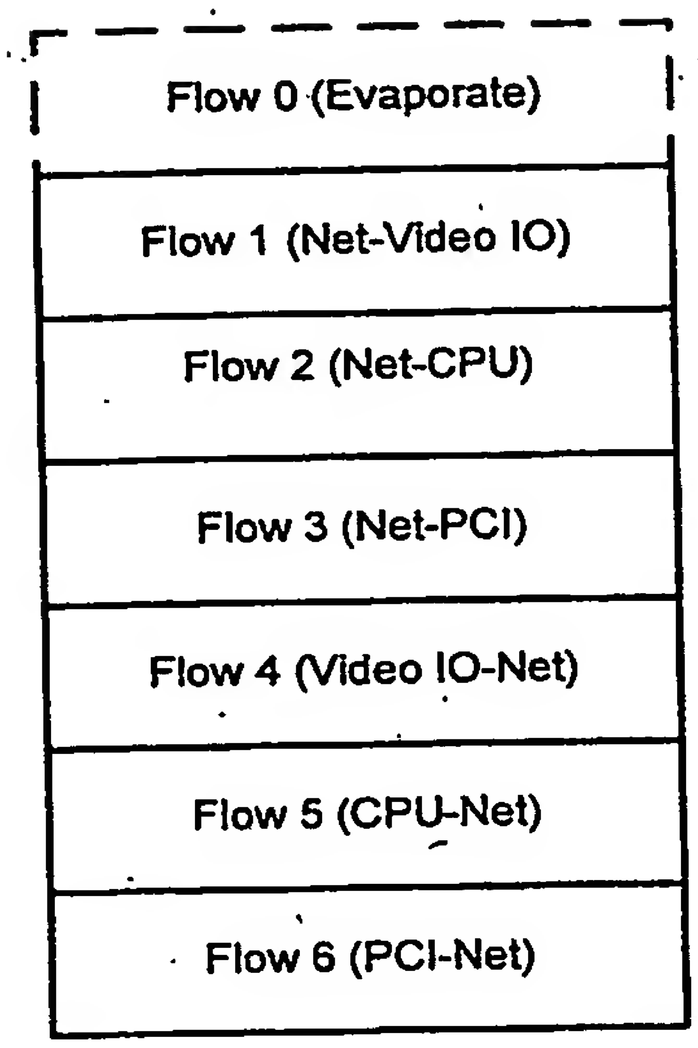

Example of the current flow assignment

FIG 5B

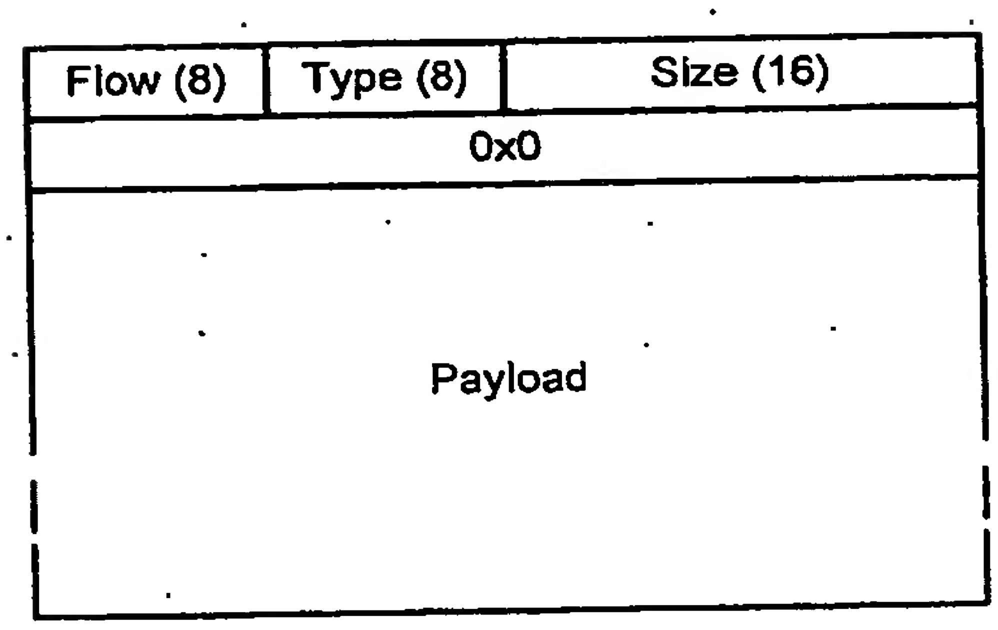

Example of a packet with a tag

FIG 5A

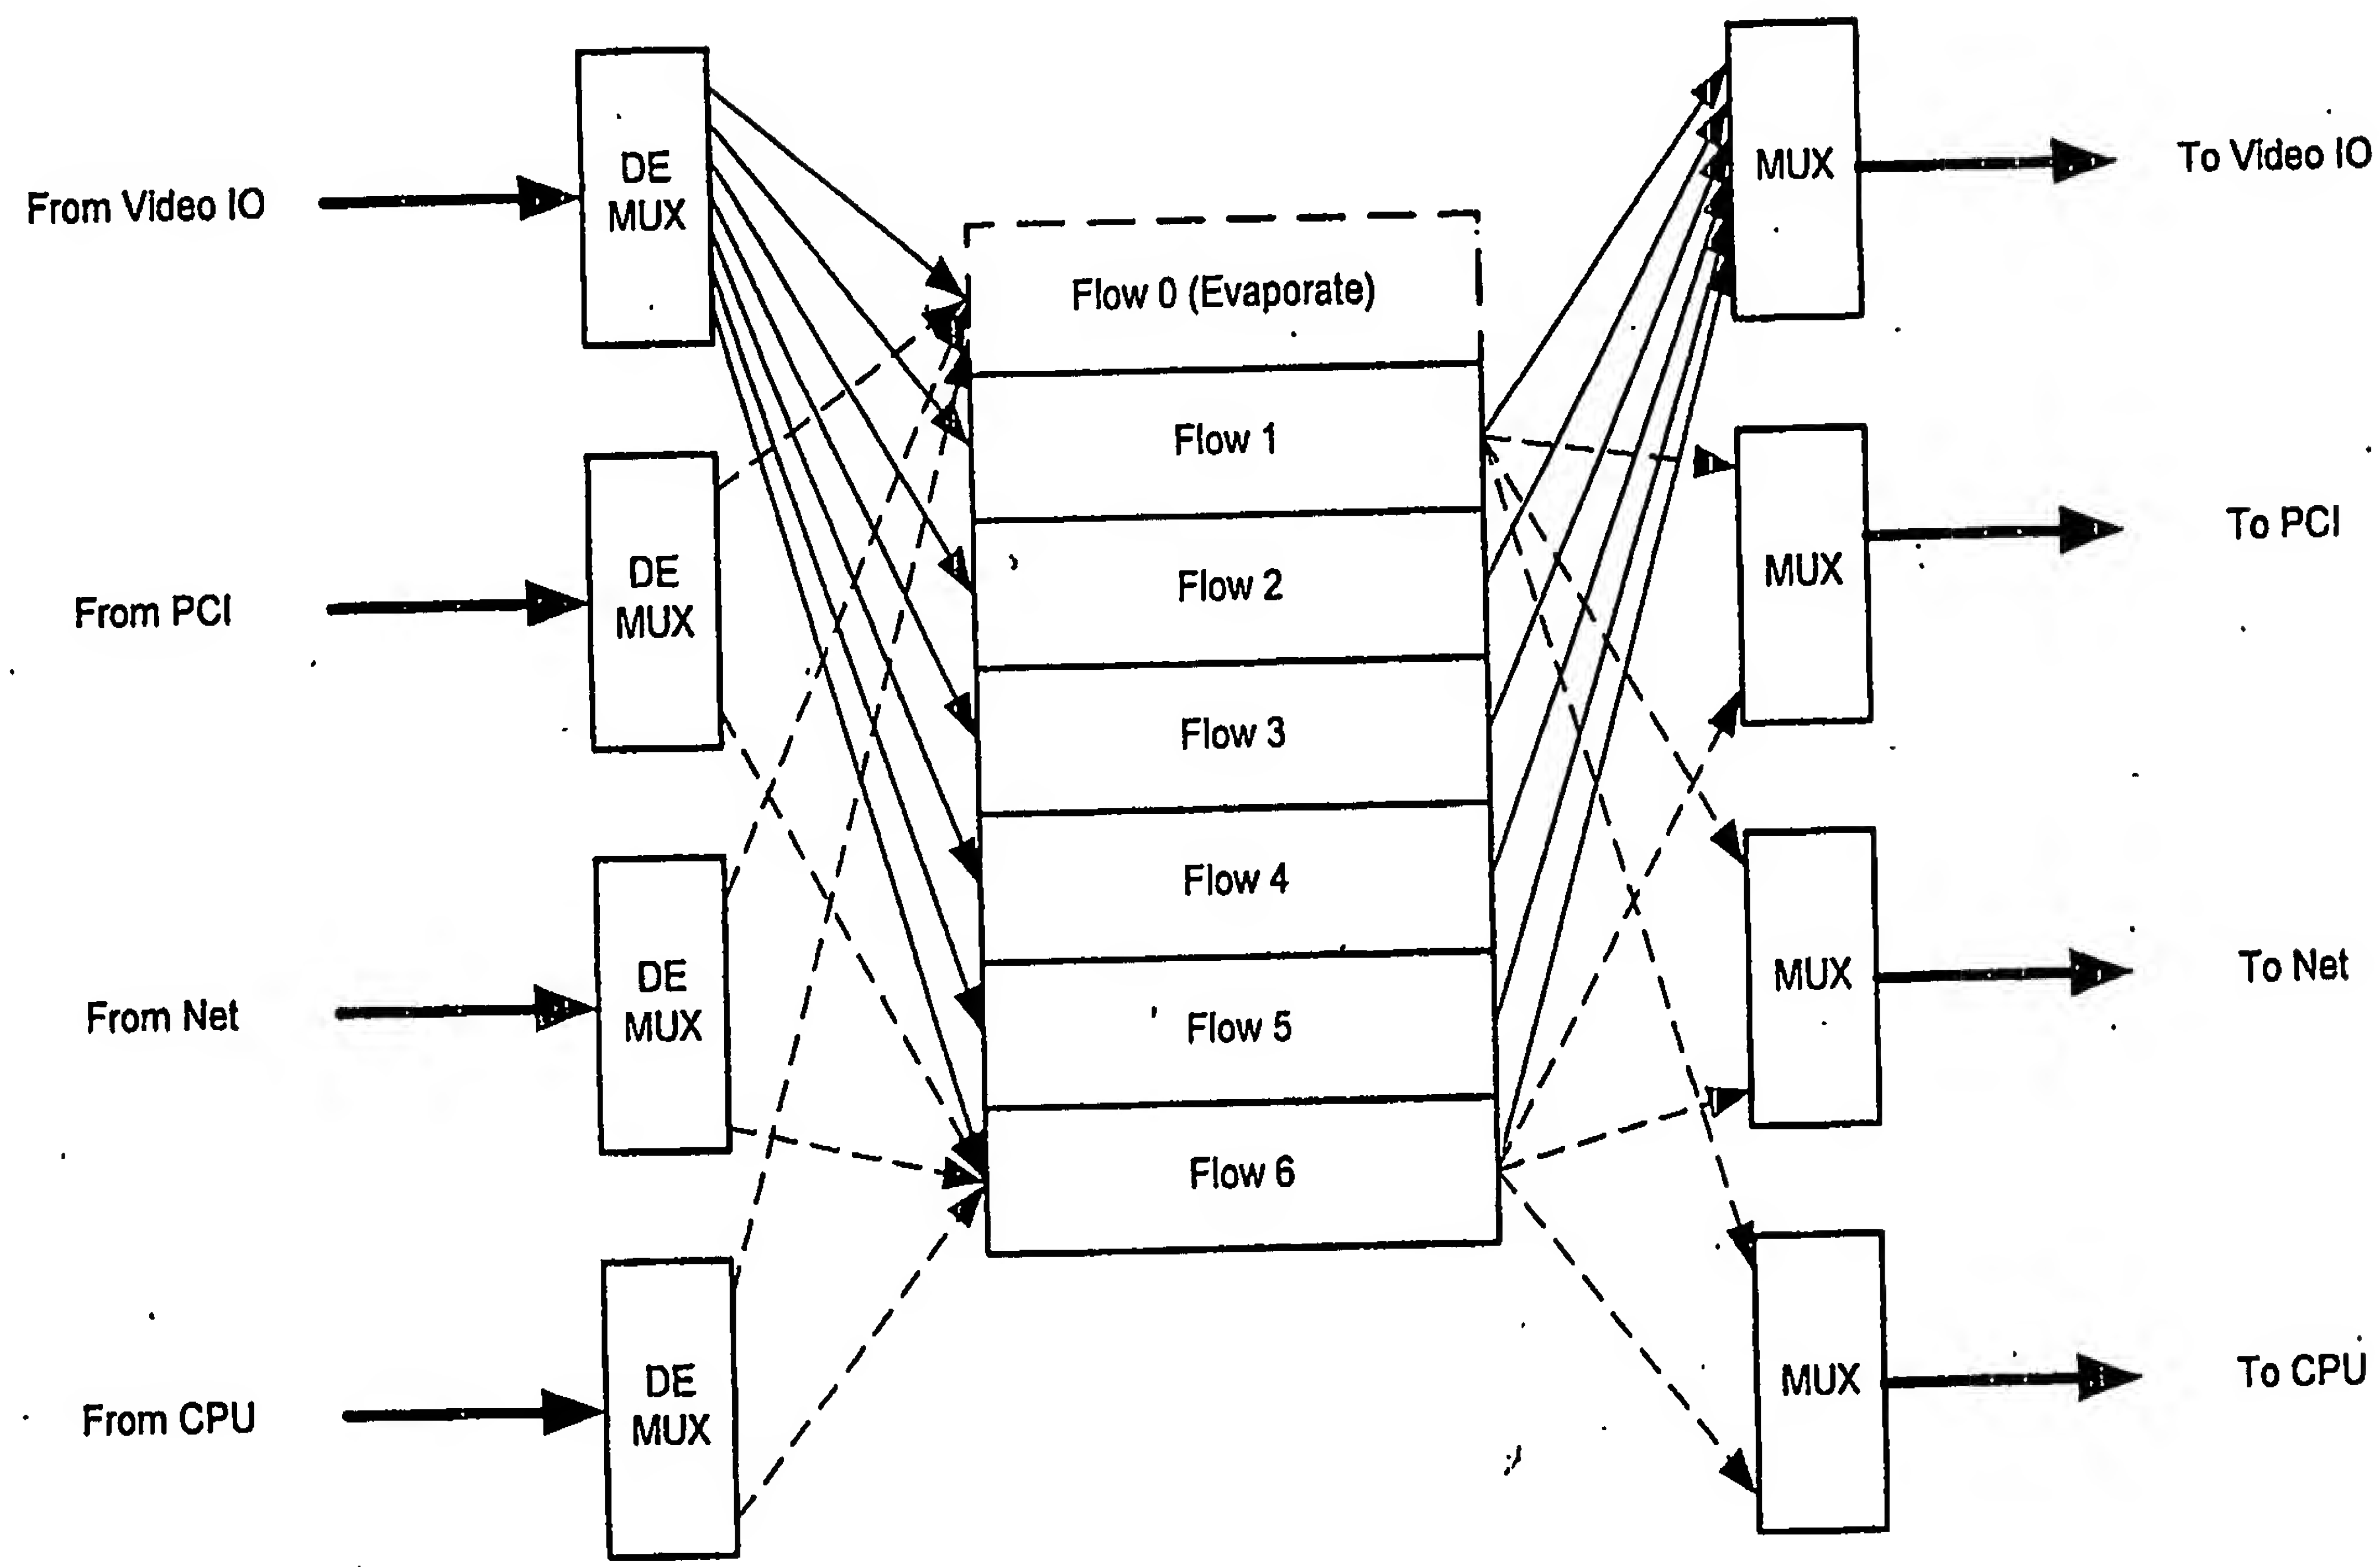
From Video IO
From PCI
From Net
From CPU
DE MUX
DE MUX
DE MUX
DE MUX
Flow 0 (Evaporate)
Flow 1
Flow 2
Flow 3
Flow 4
Flow 5
Flow 6
MUX
MUX
MUX
MUX
To Video IO
To PCI
To Net
To CPU

FIG 5C

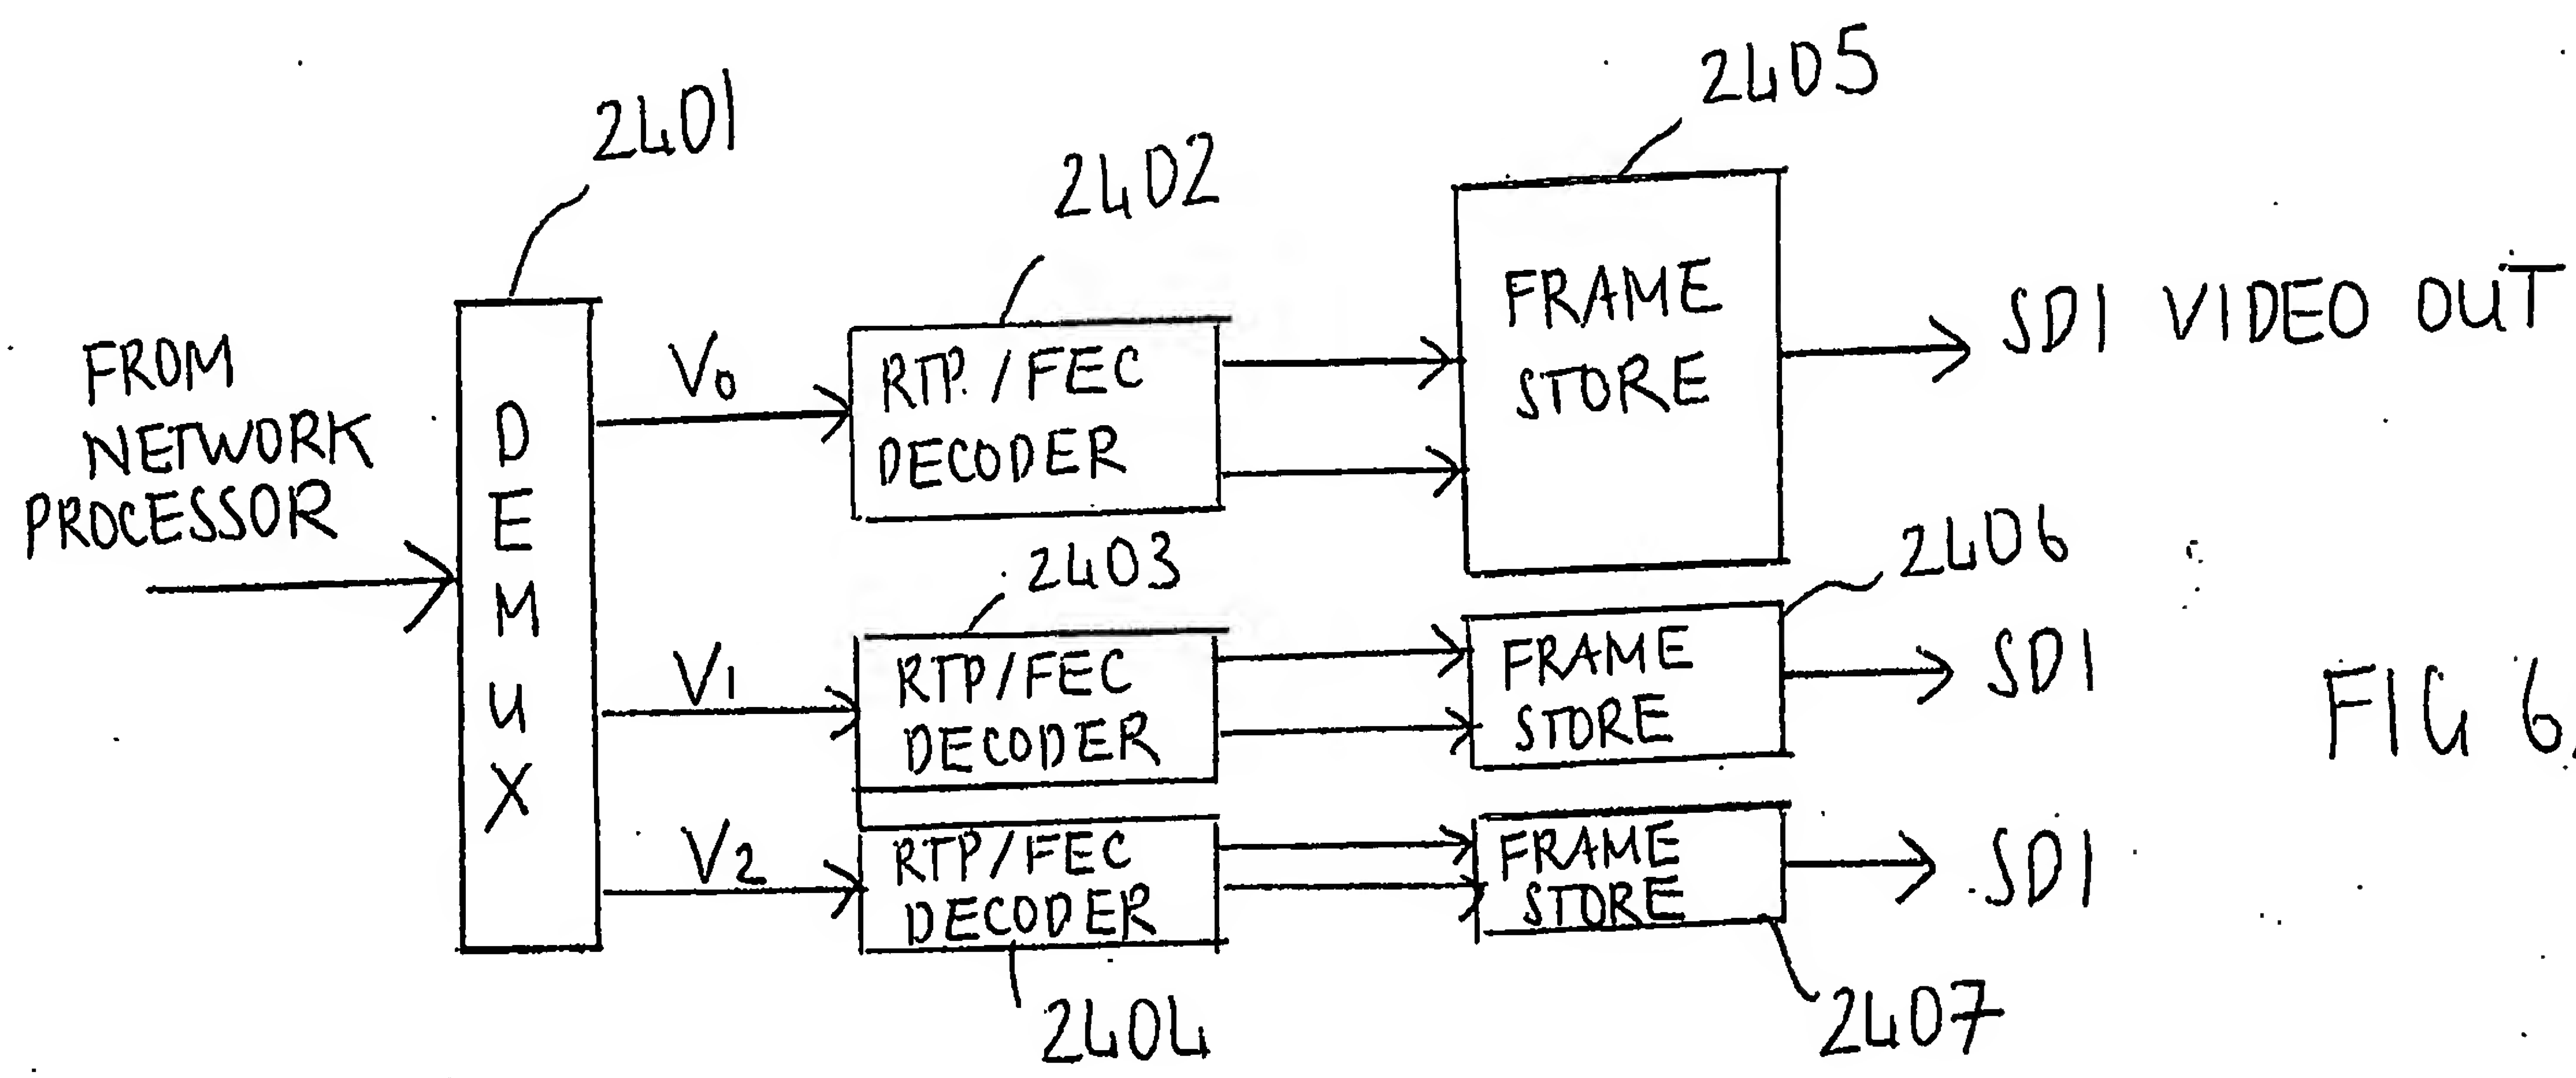
FROM NETWORK PROCESSOR
DEMUX
2401
V_0
V_1
V_2
RTP/FEC DECODER
2402
2403
2404
FRAME STORE
2405
2406
2407
SDI VIDEO OUT
SDI
SDI

FIG 6A

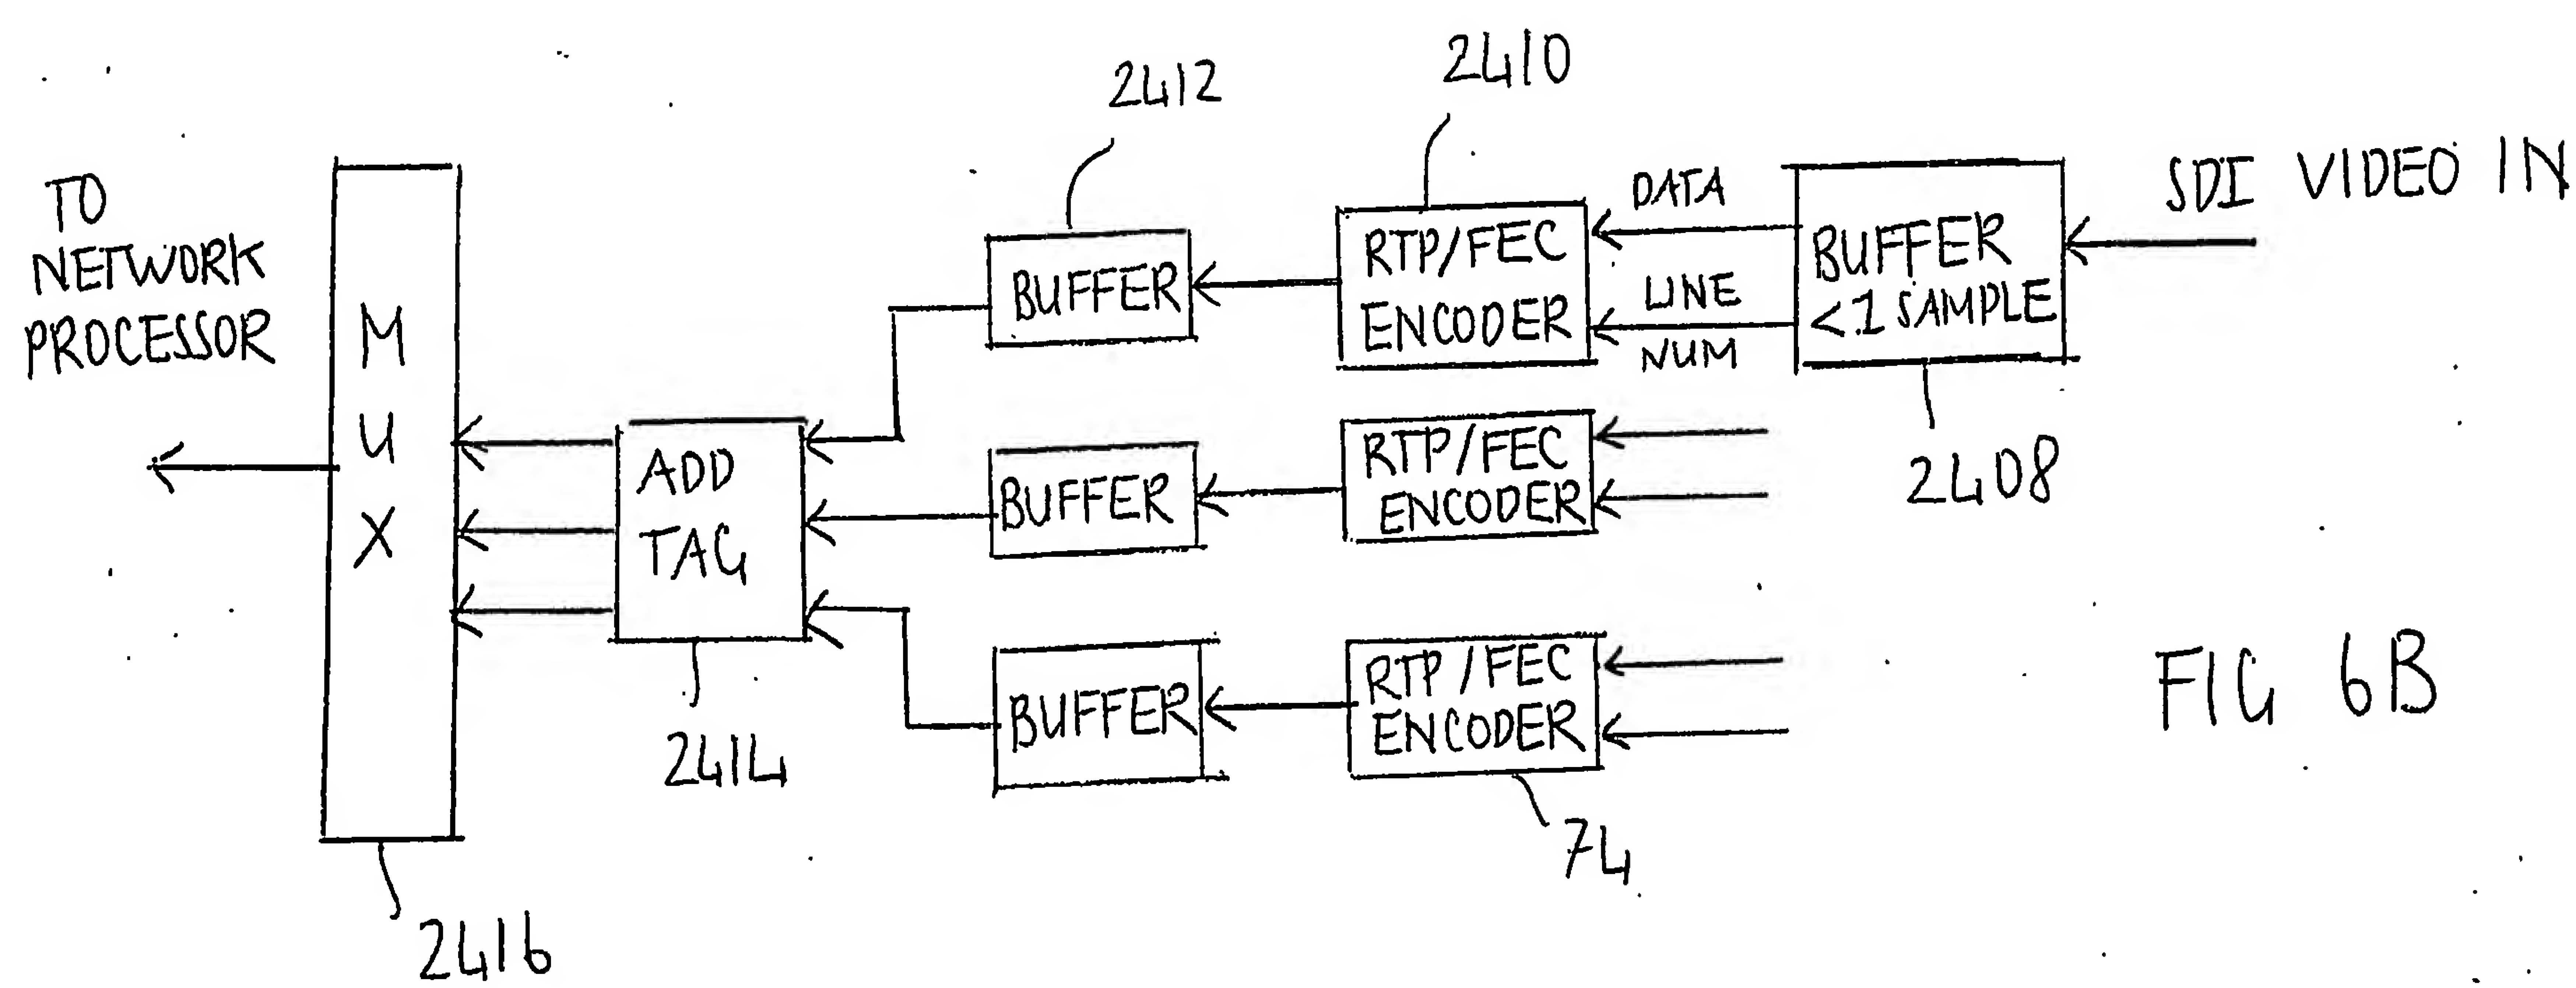
SDI VIDEO IN
BUFFER <1 SAMPLE
2408
DATA
LINE NUM
RTP/FEC ENCODER
2410
74
BUFFER
2412
ADD TAG
2414
MUX
2416
TO NETWORK PROCESSOR

FIG 6B

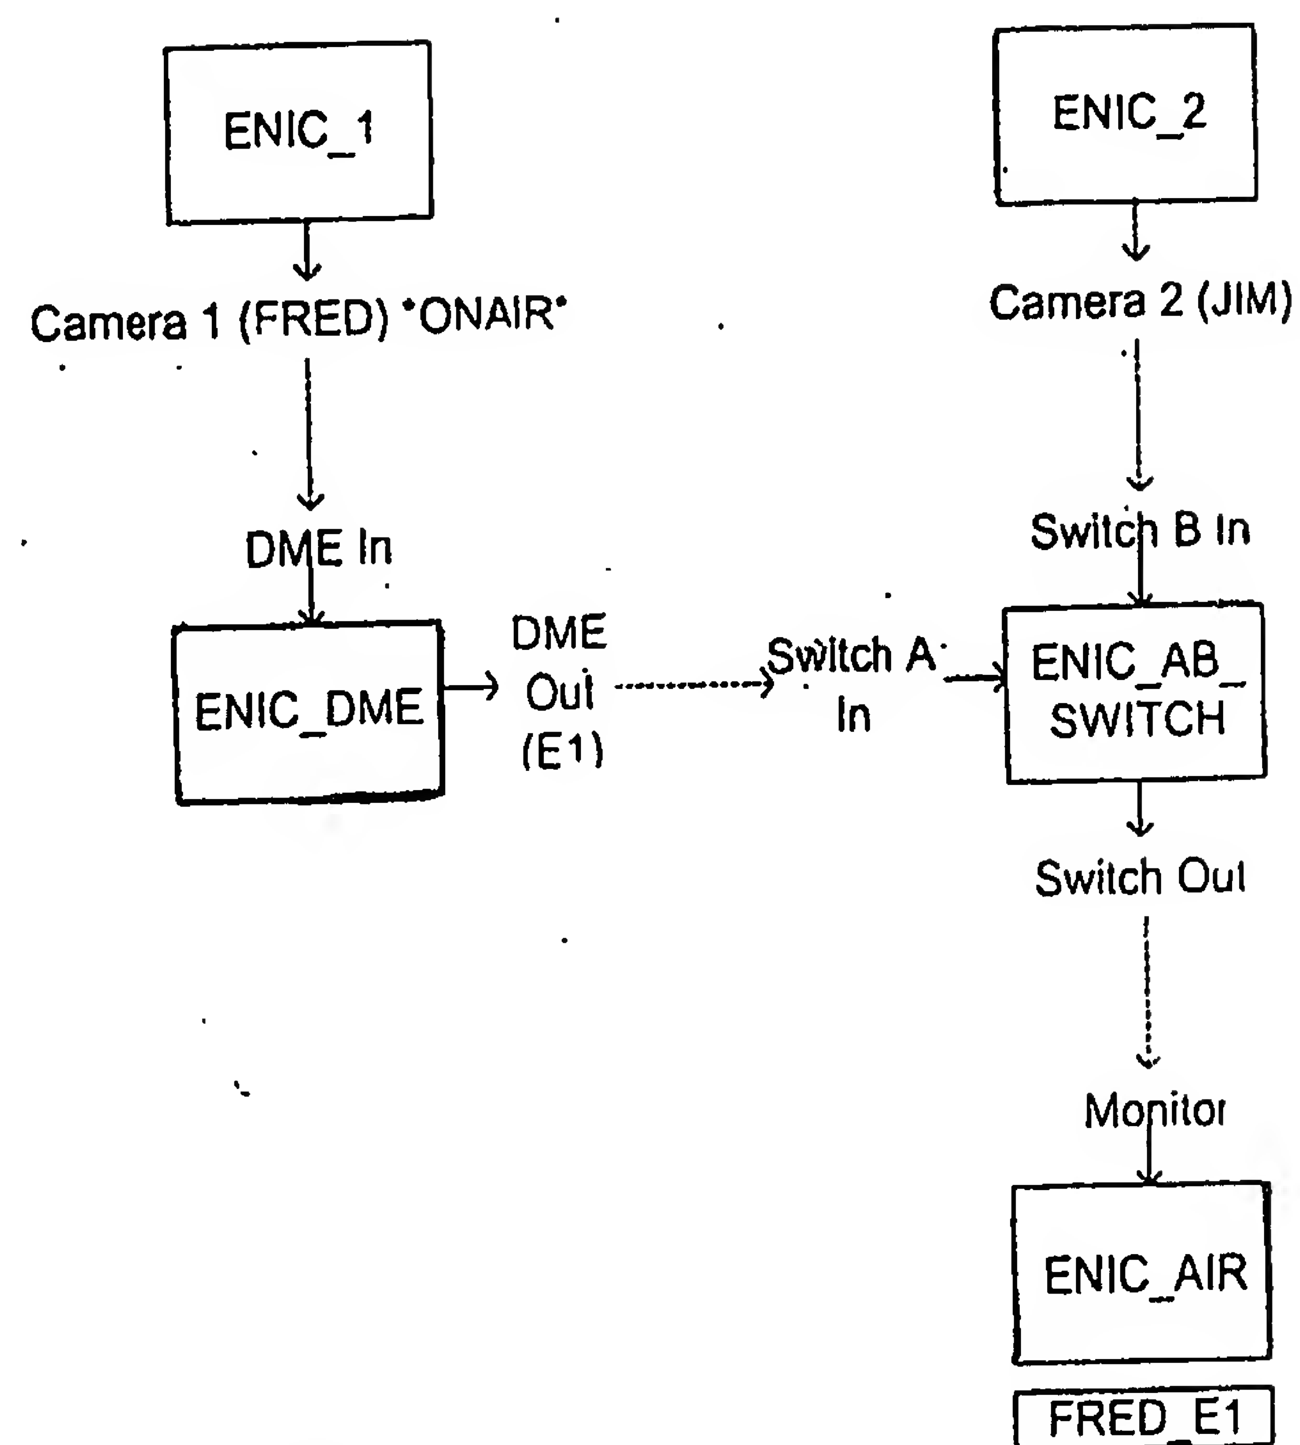
ENIC_1
Camera 1 (FRED) "ONAIR"
DME In
ENIC_DME
DME Out (E1)
ENIC_2
Camera 2 (JIM)
Switch B In
Switch A In
ENIC_AB_SWITCH
Switch Out
Monitor
ENIC_AIR
FRED_E1

Figure 7

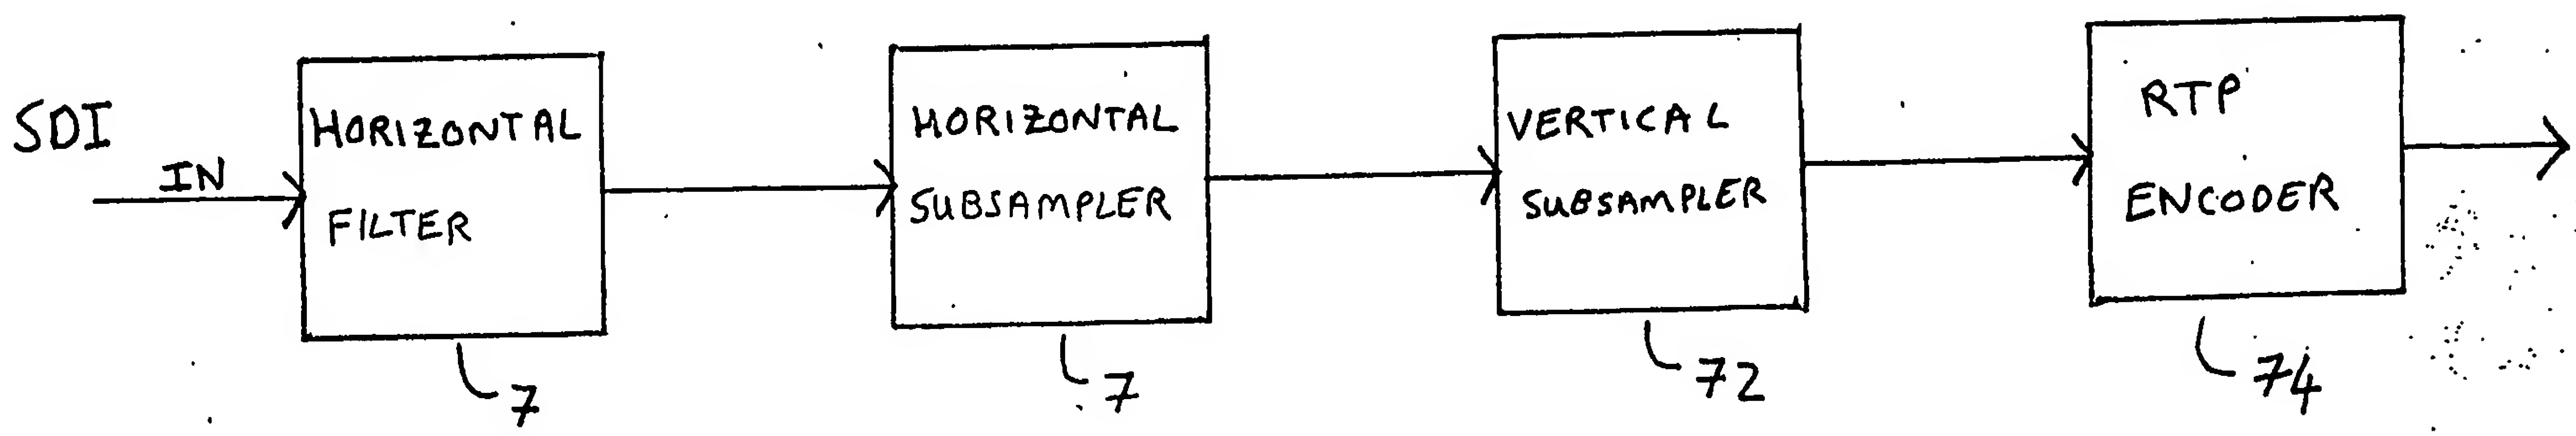
SDI
IN
HORIZONTAL FILTER
7
HORIZONTAL SUBSAMPLER
7
VERTICAL SUBSAMPLER
72
RTP ENCODER
74

Figure 8

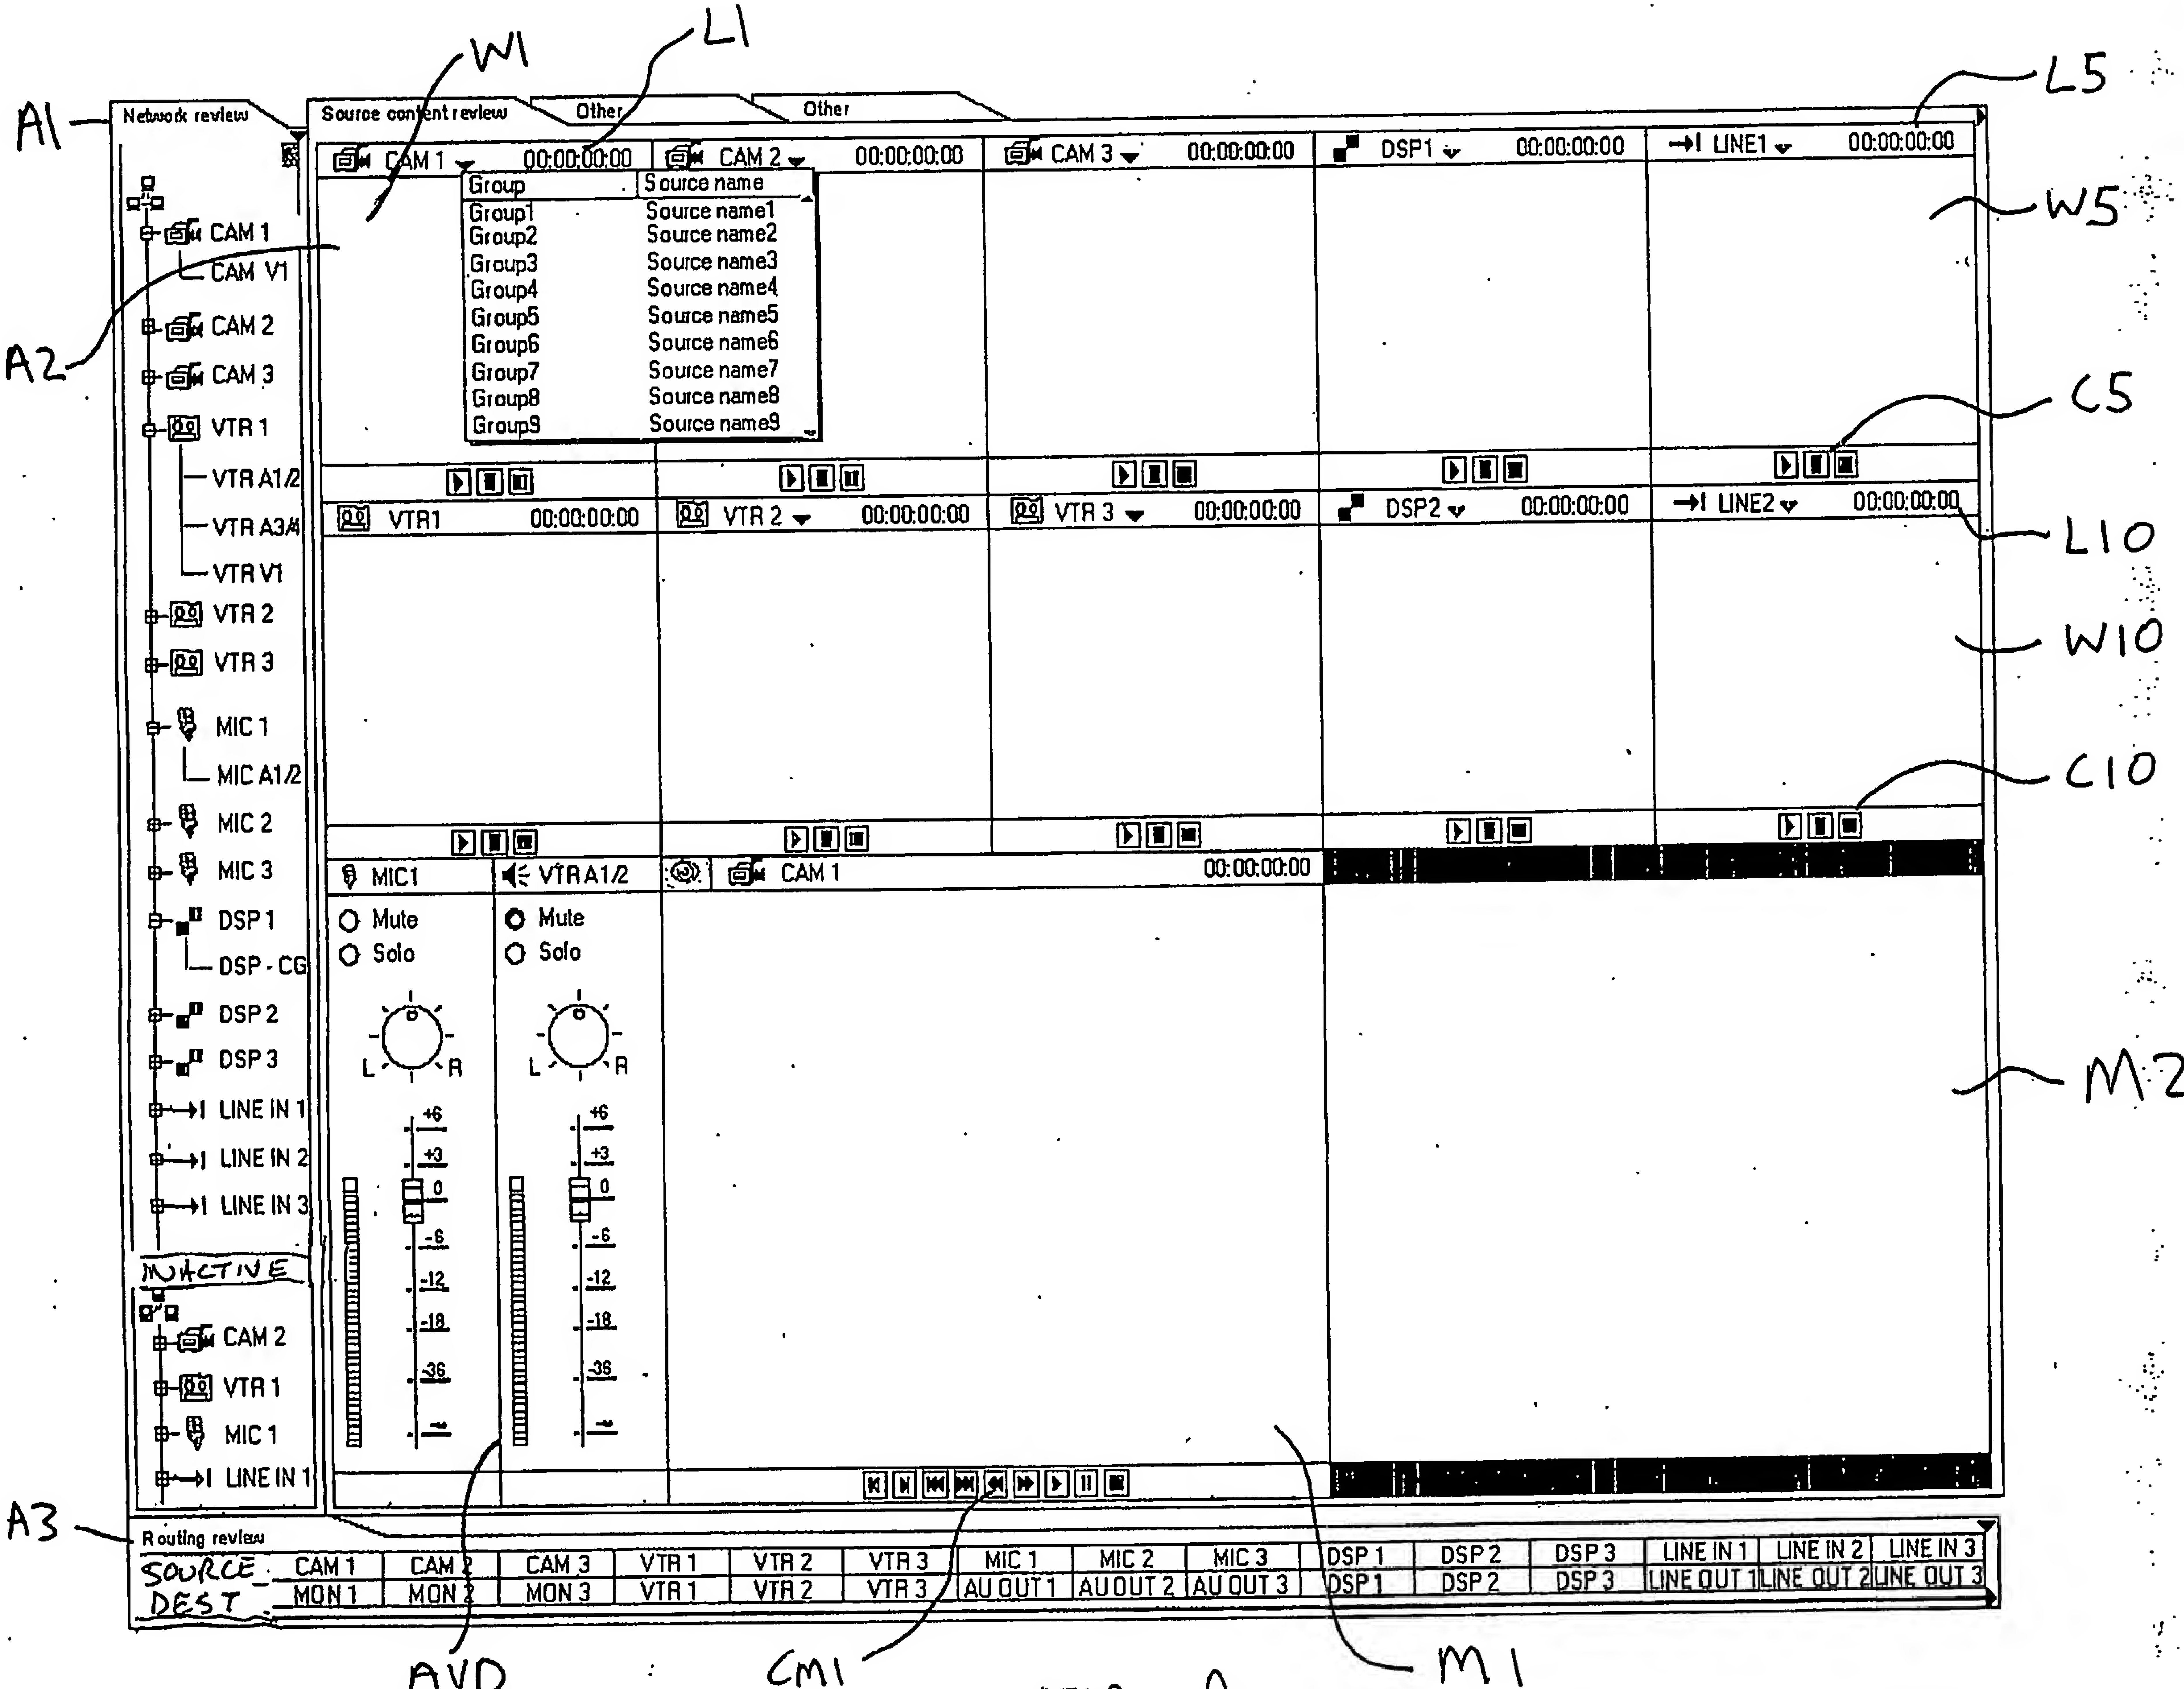
A1
A2
A3
W1
L1
L5
W5
C5
L10
W10
C10
M2
M1
CM1
AVD
Network review
Source content review
Other
Other
Routing review
CAM 1
CAM V1
CAM 2
CAM 3
VTR 1
VTR A1/2
VTR A3/4
VTR V1
VTR 2
VTR 3
MIC 1
MIC A1/2
MIC 2
MIC 3
DSP 1
DSP - CG
DSP 2
DSP 3
LINE IN 1
LINE IN 2
LINE IN 3
INACTIVE
CAM 2
VTR 1
MIC 1
LINE IN 1
CAM 1 00:00:00:00
CAM 2 00:00:00:00
CAM 3 00:00:00:00
DSP1 00:00:00:00
LINE1 00:00:00:00
Group
Group1
Group2
Group3
Group4
Group5
Group6
Group7
Group8
Group9
Source name
Source name1
Source name2
Source name3
Source name4
Source name5
Source name6
Source name7
Source name8
Source name9
VTR1 00:00:00:00
VTR 2 00:00:00:00
VTR 3 00:00:00:00
DSP2 00:00:00:00
LINE2 00:00:00:00
MIC1
VTR A1/2
CAM 1 00:00:00:00
Mute
Solo
L R
+6
+3
0
-6
-12
-18
-36
SOURCE
DEST
CAM 1 | MON 1
CAM 2 | MON 2
CAM 3 | MON 3
VTR 1 | VTR 1
VTR 2 | VTR 2
VTR 3 | VTR 3
MIC 1 | AU OUT 1
MIC 2 | AU OUT 2
MIC 3 | AU OUT 3
DSP 1 | DSP 1
DSP 2 | DSP 2
DSP 3 | DSP 3
LINE IN 1 | LINE OUT 1
LINE IN 2 | LINE OUT 2
LINE IN 3 | LINE OUT 3

FIG 9

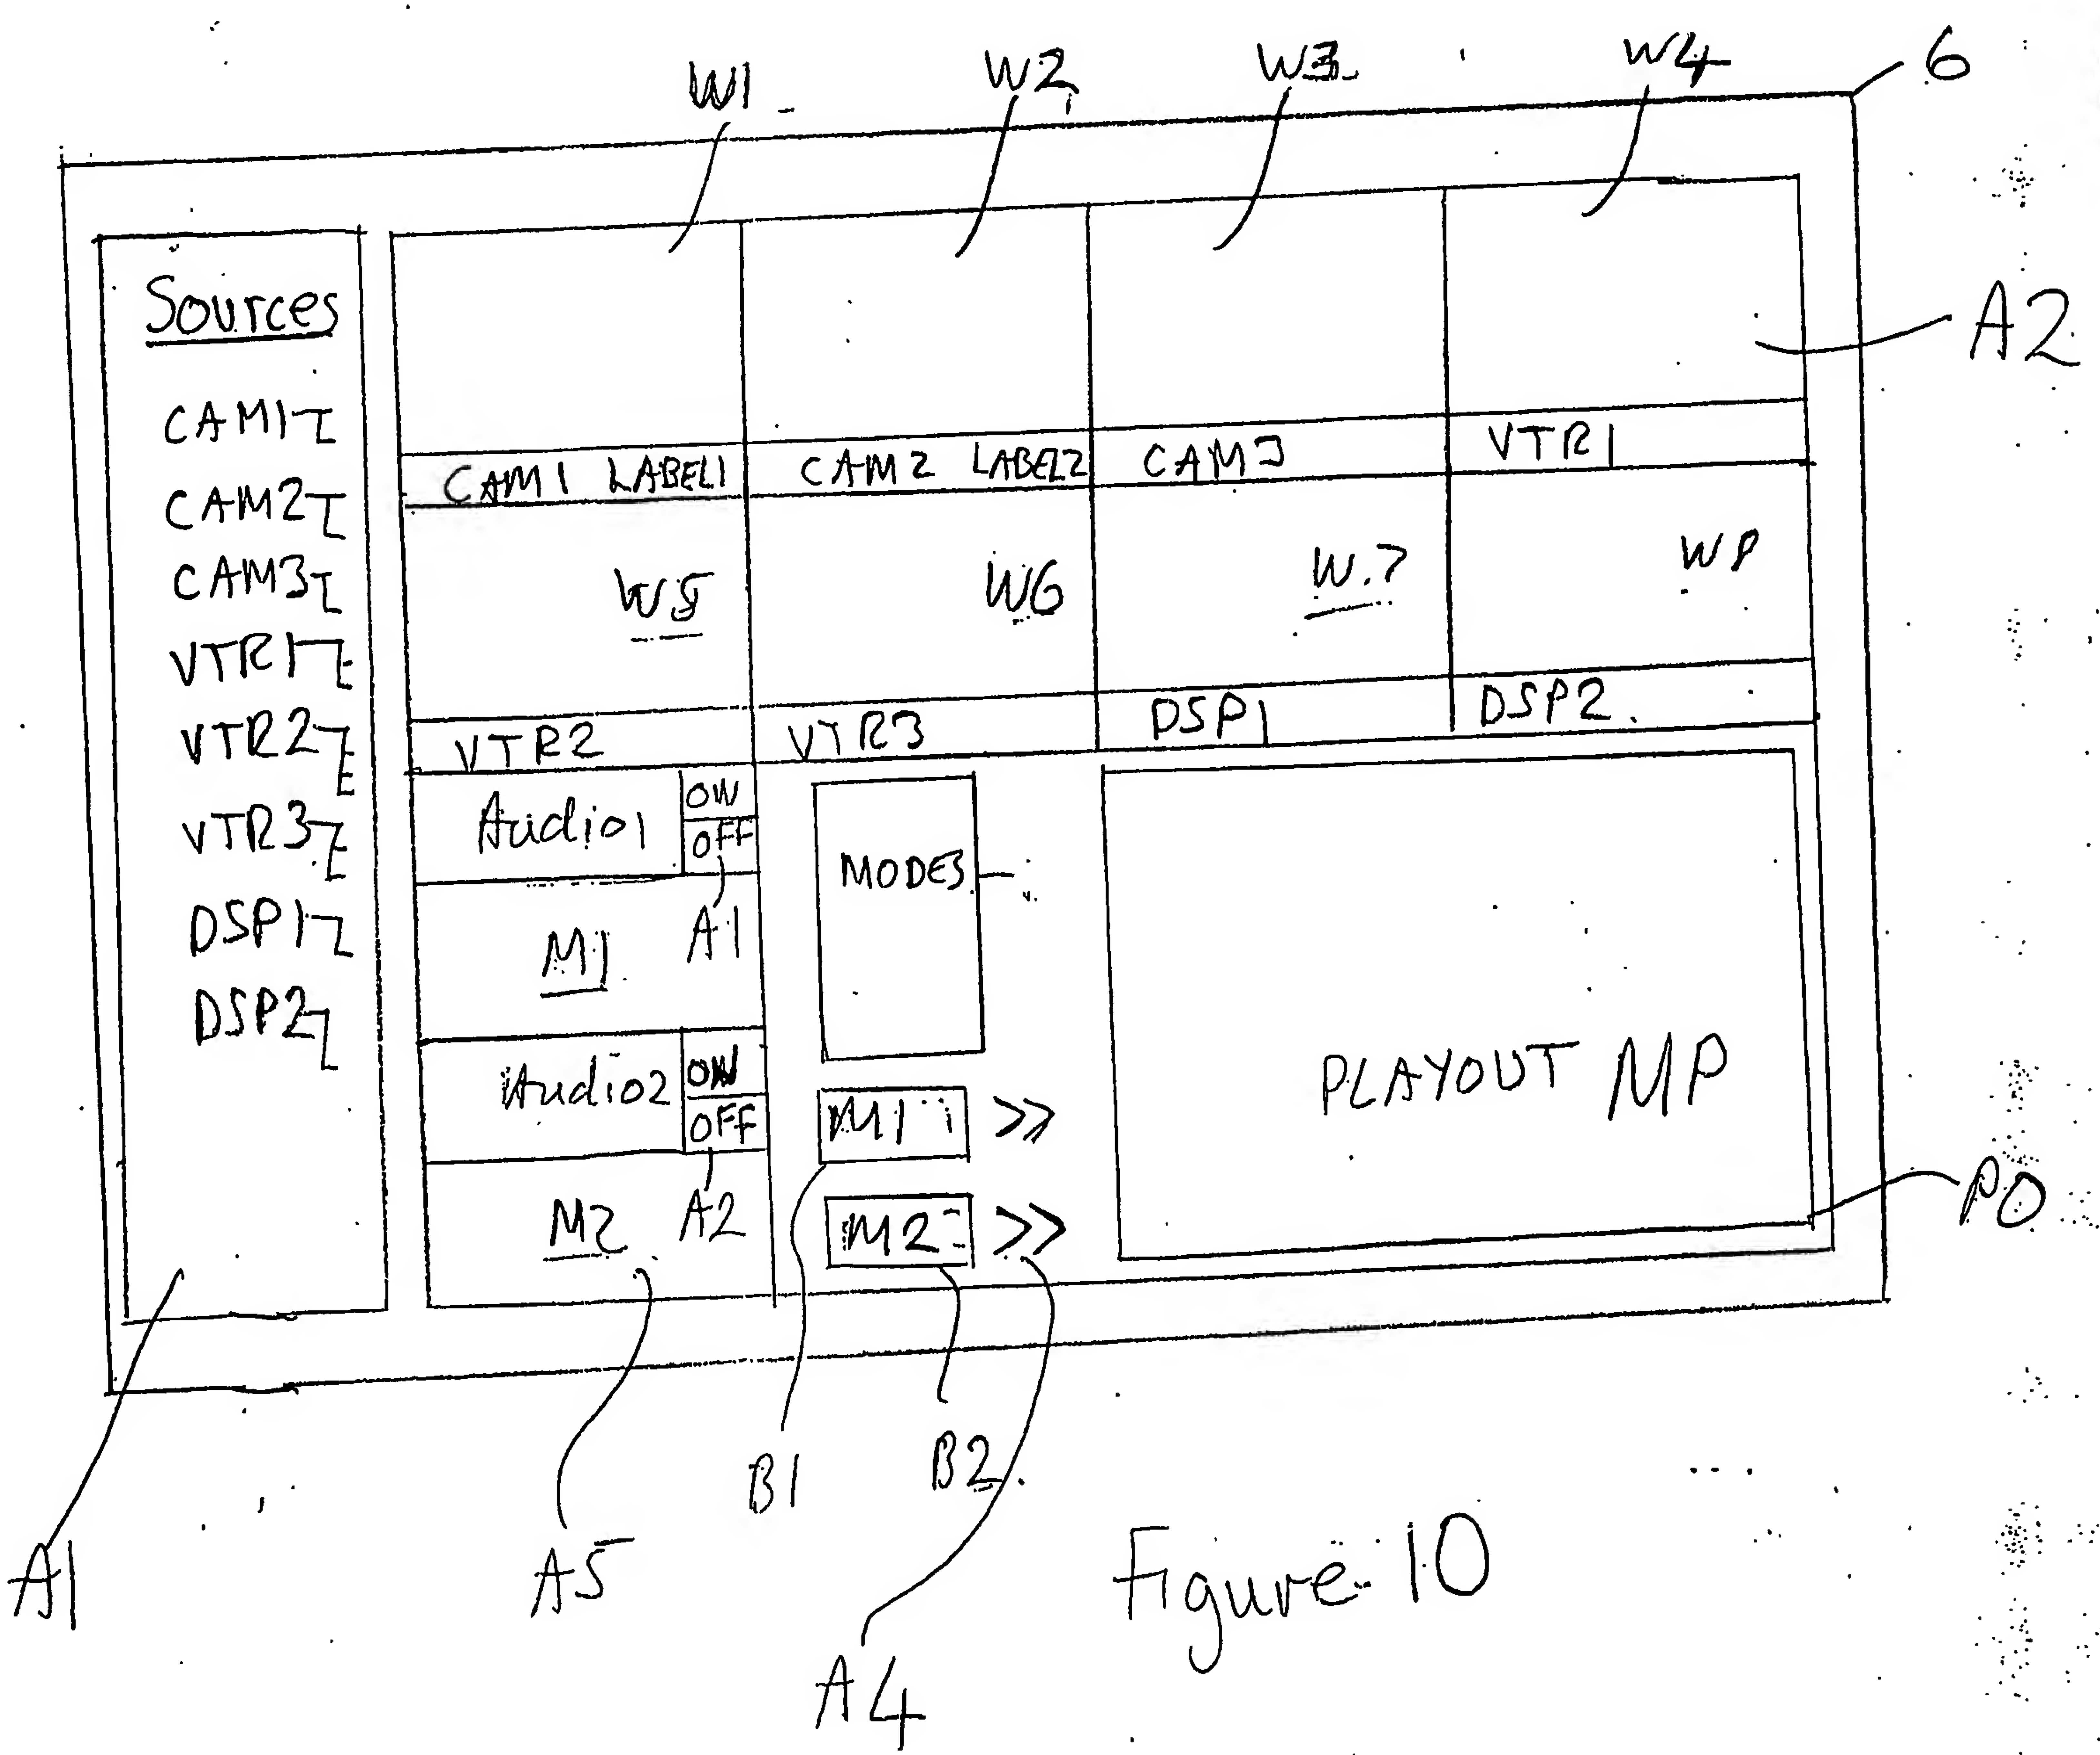
W1
W2
W3
W4
6
A2
Sources
CAM1
CAM2
CAM3
VTR1
VTR2
VTR3
DSP1
DSP2
CAM1 LABEL1
CAM2 LABEL2
CAM3
VTR1
W5
W6
W7
W8
VTR2
VTR3
DSP1
DSP2
Audio1
ON
OFF
M1
A1
Audio2
ON
OFF
M2
A2
MODES
M1
M2
PLAYOUT MP
PO
A1
A5
B1
B2
A4

Figure 10

10

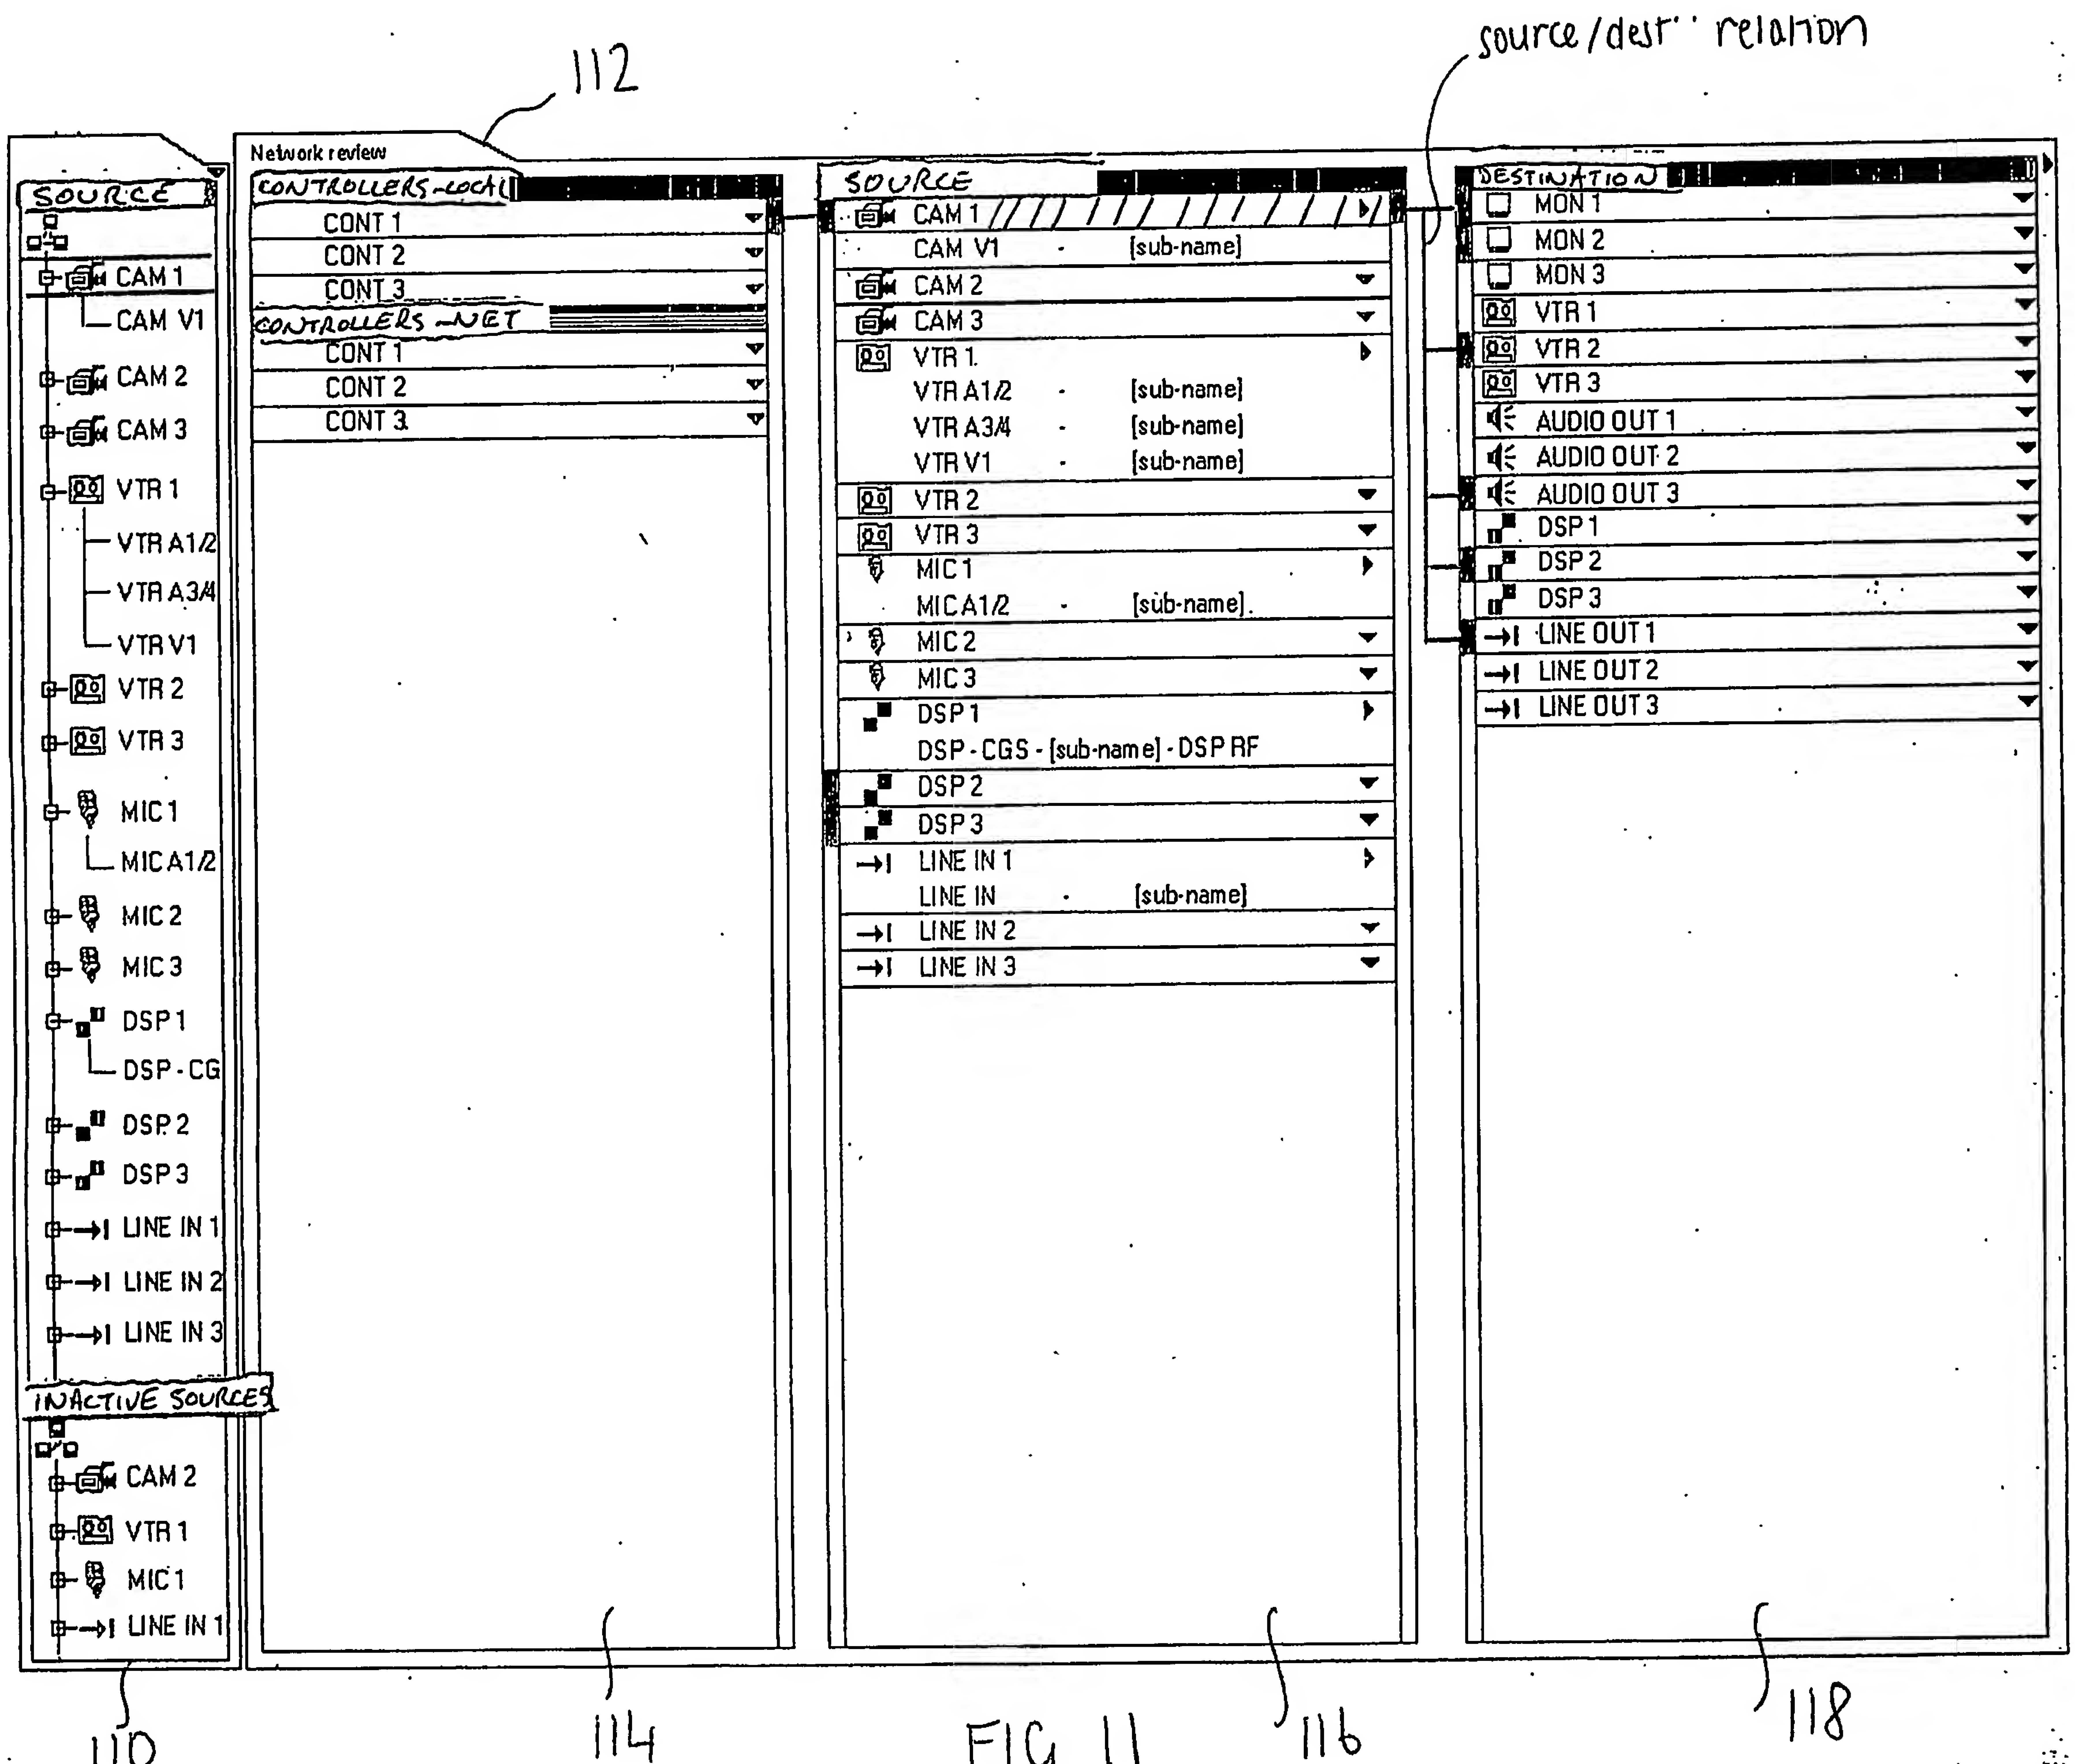
source/dest' relation
112
SOURCE
CAM 1
CAM V1
CAM 2
CAM 3
VTR 1
VTR A1/2
VTR A3/4
VTR V1
VTR 2
VTR 3
MIC 1
MIC A1/2
MIC 2
MIC 3
DSP 1
DSP-CG
DSP 2
DSP 3
LINE IN 1
LINE IN 2
LINE IN 3
INACTIVE SOURCES
CAM 2
VTR 1
MIC 1
LINE IN 1
Network review
CONTROLLERS-LOCAL
CONT 1
CONT 2
CONT 3
CONTROLLERS-NET
CONT 1
CONT 2
CONT 3
SOURCE
CAM 1
CAM V1 - [sub-name]
CAM 2
CAM 3
VTR 1
VTR A1/2 - [sub-name]
VTR A3/4 - [sub-name]
VTR V1 - [sub-name]
VTR 2
VTR 3
MIC 1
MIC A1/2 - [sub-name]
MIC 2
MIC 3
DSP 1
DSP-CGS-[sub-name]-DSP RF
DSP 2
DSP 3
LINE IN 1
LINE IN - [sub-name]
LINE IN 2
LINE IN 3
DESTINATION
MON 1
MON 2
MON 3
VTR 1
VTR 2
VTR 3
AUDIO OUT 1
AUDIO OUT 2
AUDIO OUT 3
DSP 1
DSP 2
DSP 3
LINE OUT 1
LINE OUT 2
LINE OUT 3
110
114
116
118
FIG 11

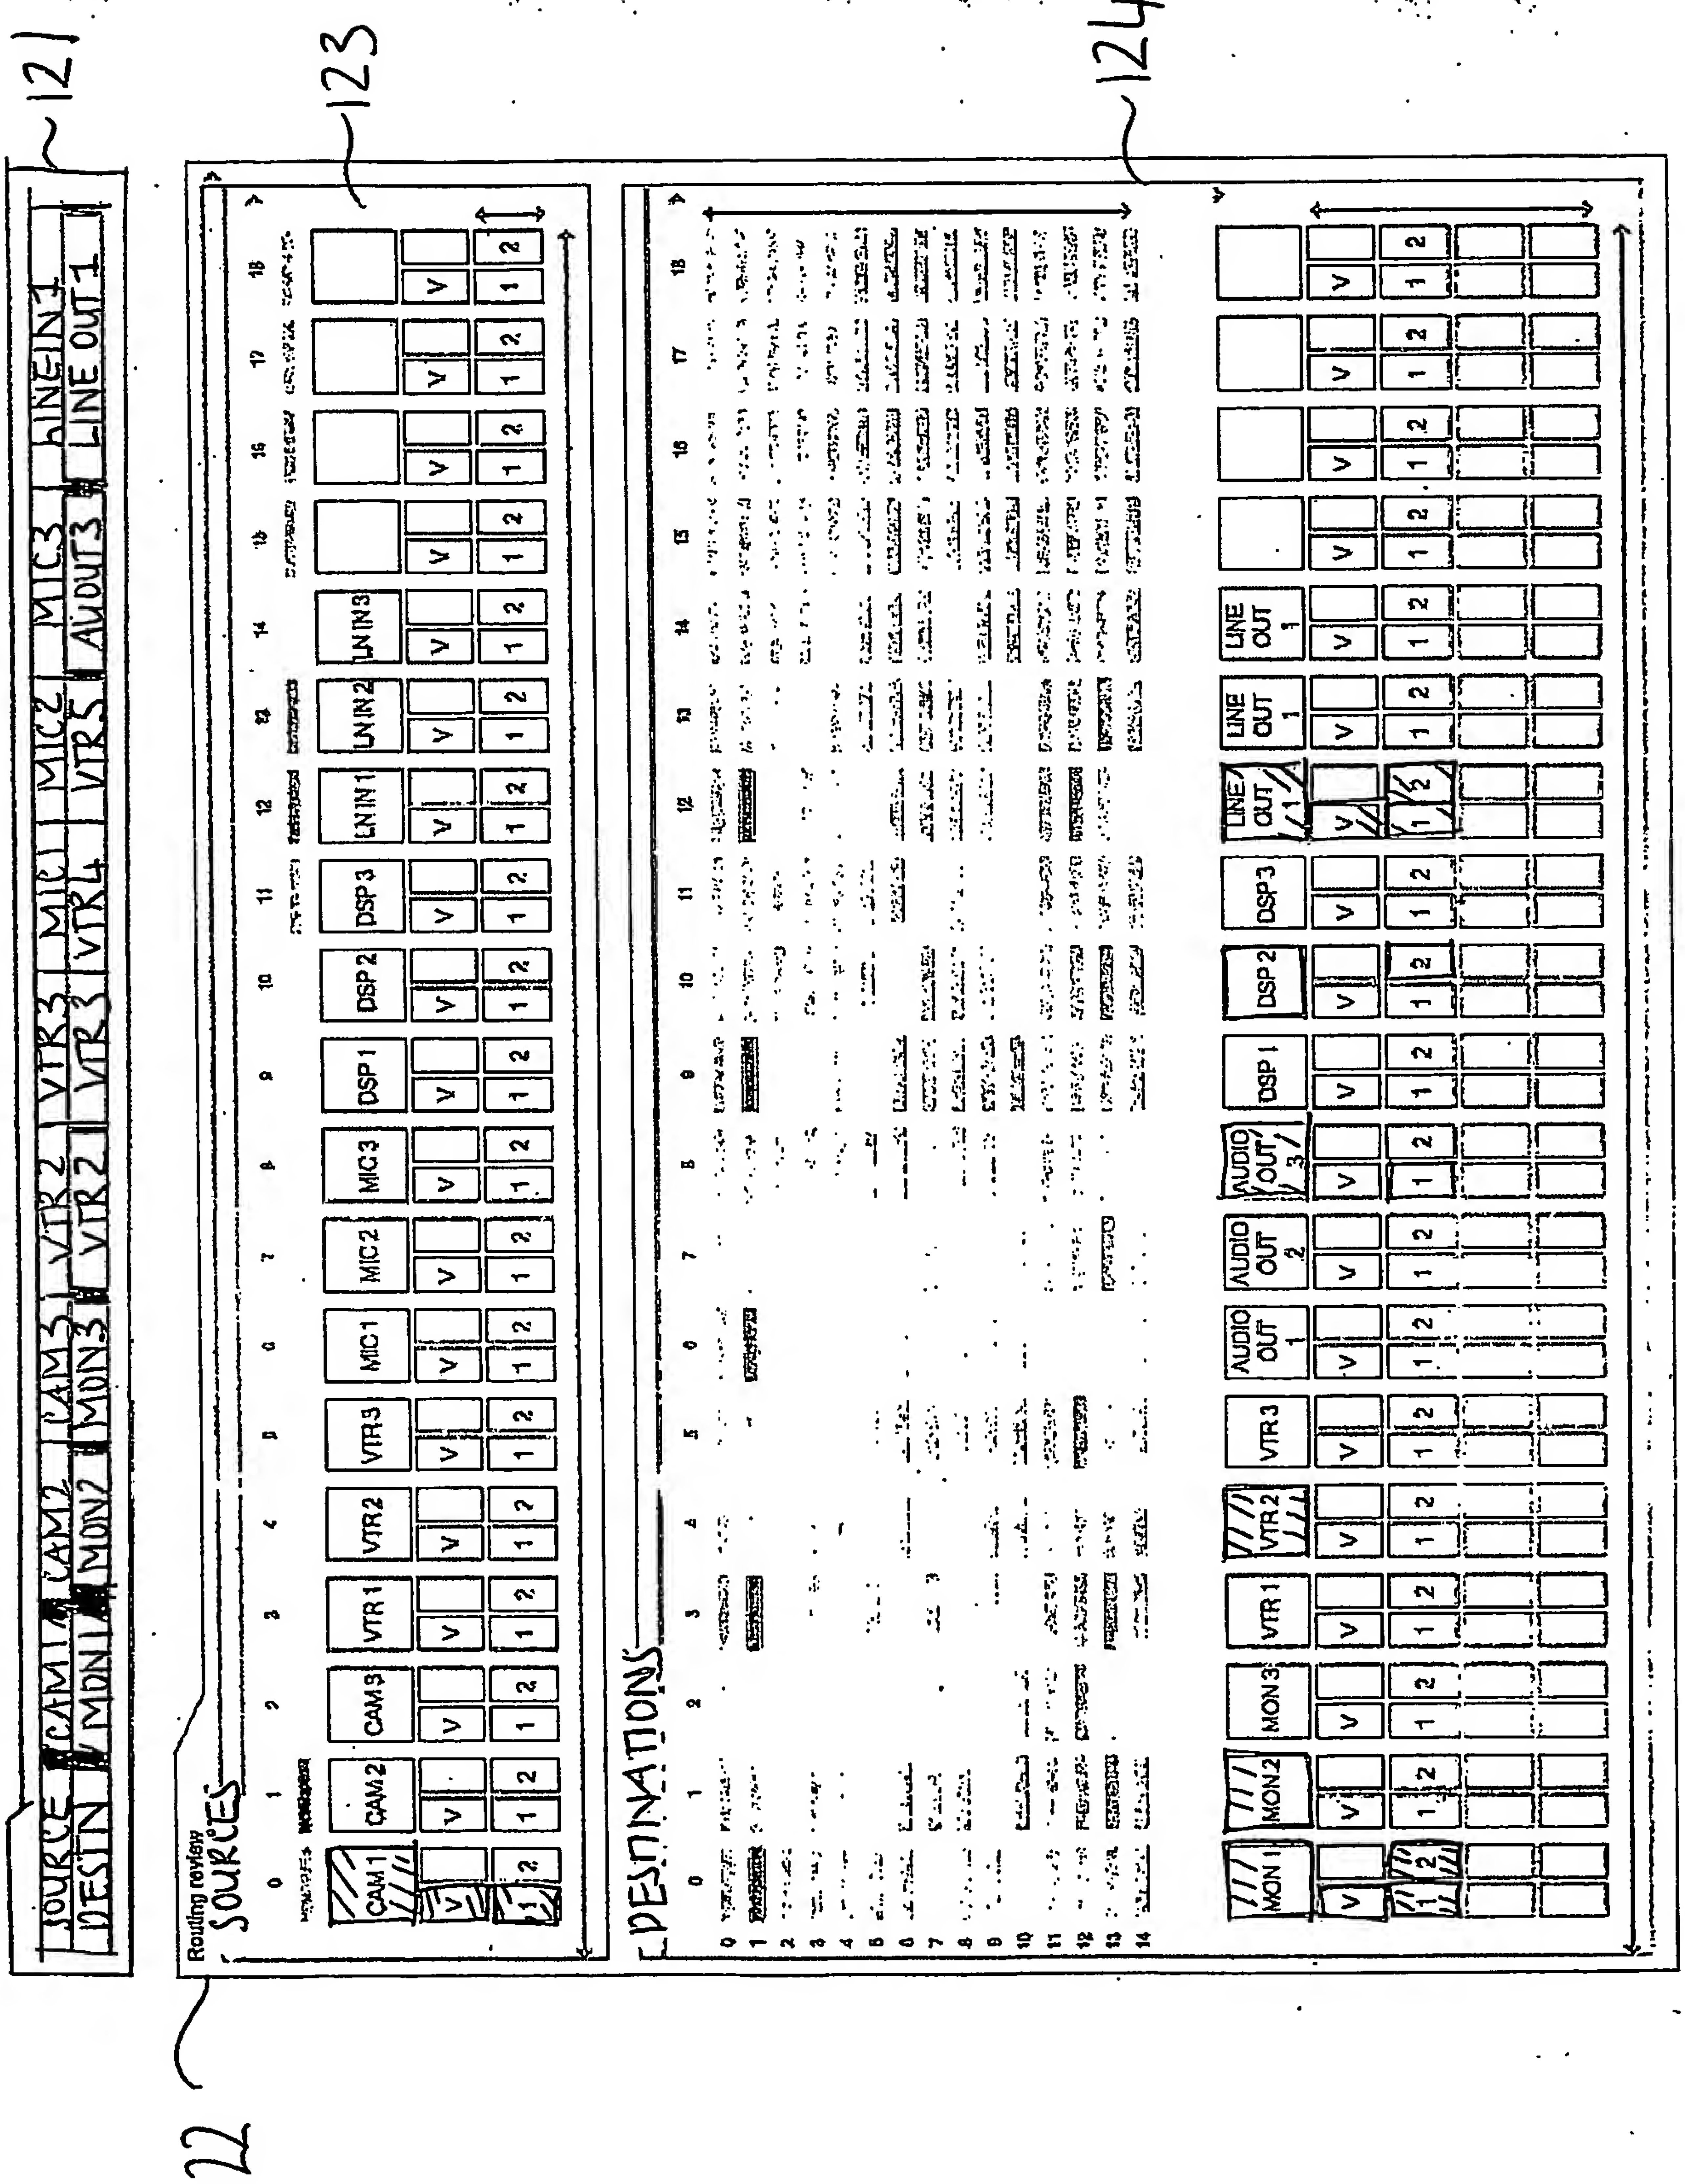
121
SOURCE
DESTN
CAM1
CAM2
CAM3
VTR2
VTR3
MIC1
MIC2
MIC3
LINEIN1
MON1
MON2
MON3
VTR2
VTR3
VTR4
VTR5
AUDOUT3
LINE OUT1
123
124
Routing review
SOURCES
DESTINATIONS
CAM1
CAM2
CAM3
VTR1
VTR2
VTR3
MIC1
MIC2
MIC3
DSP1
DSP2
DSP3
LN IN1
LN IN2
LN IN3
MON1
MON2
MON3
AUDIO OUT 1
AUDIO OUT 2
AUDIO OUT 3
LINE OUT 1
122

FIG 12

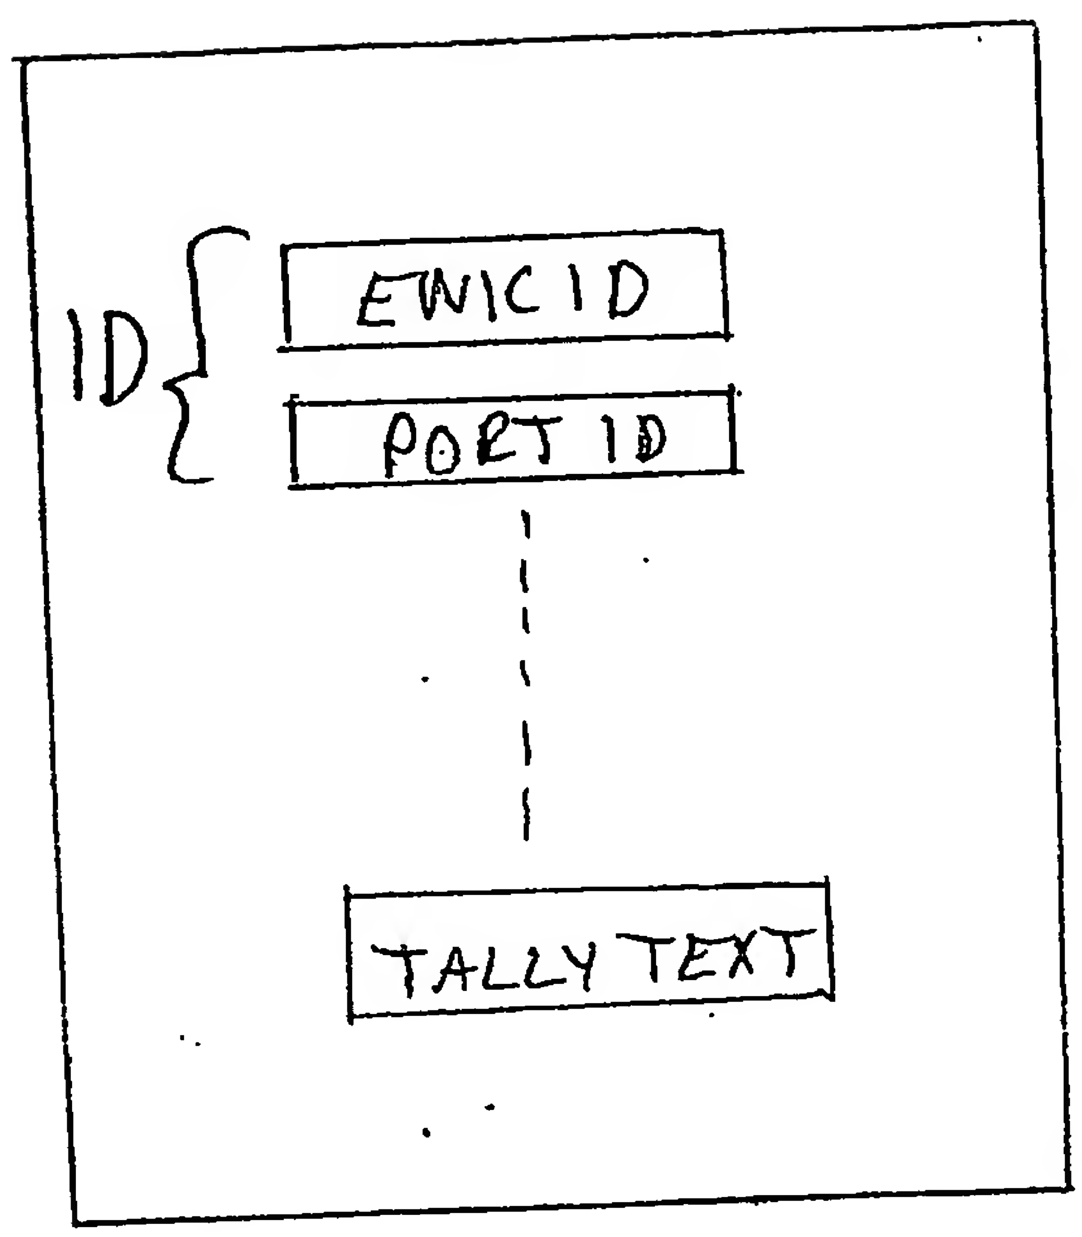

Figure 13

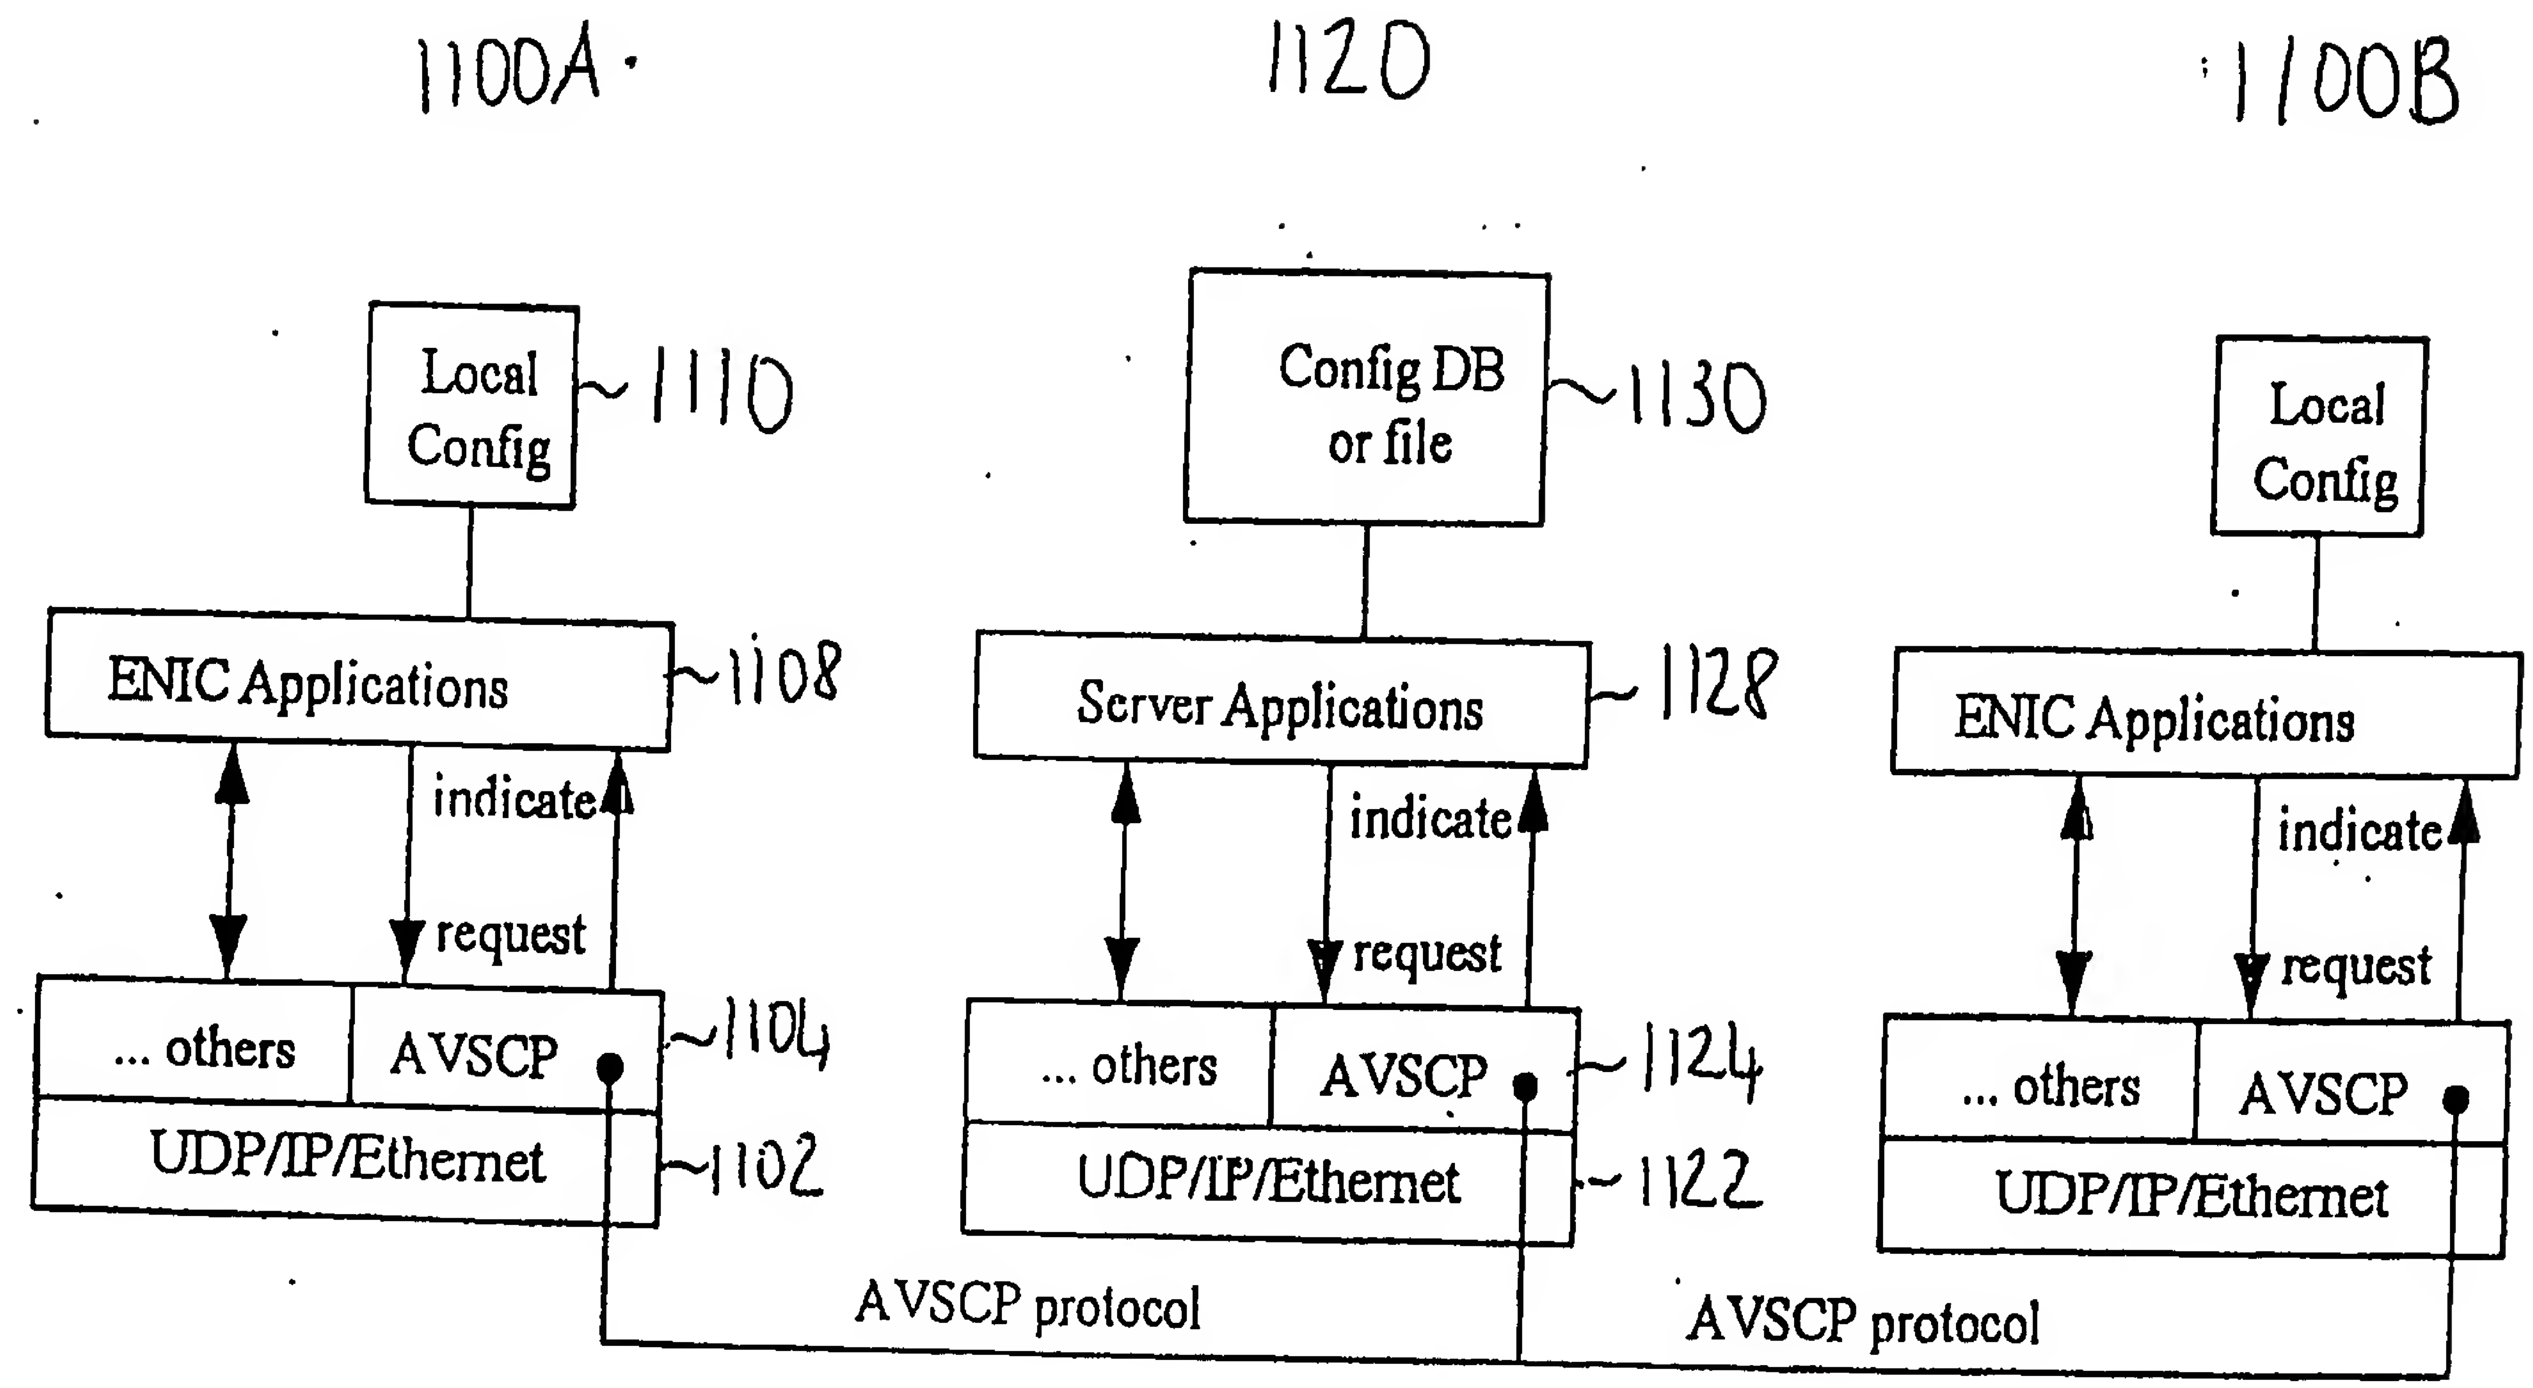
1100A
1120
1100B
Local Config
1110
Config DB or file
1130
Local Config
ENIC Applications
1108
Server Applications
1128
ENIC Applications
indicate
request
indicate
request
indicate
request
... others
AVSCP
1104
... others
AVSCP
1124
... others
AVSCP
UDP/IP/Ethernet
1102
UDP/IP/Ethernet
1122
UDP/IP/Ethernet
AVSCP protocol
AVSCP protocol

FIG 14

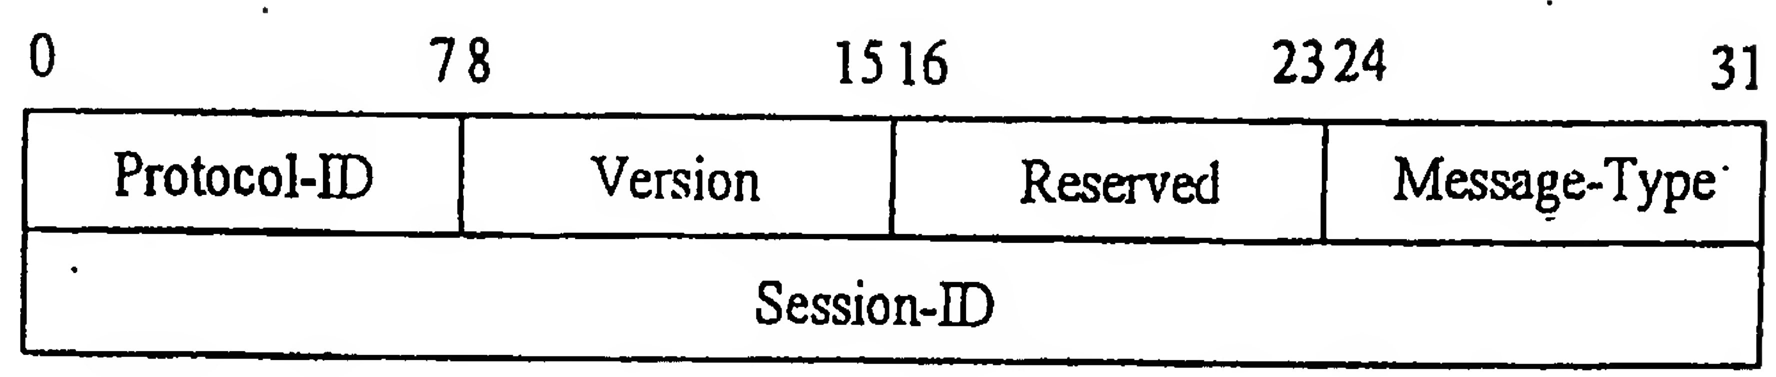
0
7 8
15 16
23 24
31
Protocol-ID
Version
Reserved
Message-Type
Session-ID

FIG 15